An Interop Publication

CONNEXIONS

The Interoperability Report

June 1996 Volume 10, No. 6

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop Company, a division of SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California, 94404–1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: `connexions@interop.com`

Subscription hotline: 1-800-575-5717
or +1 610-892-1959

ISSN 0894–5926

From the Editor

In our June 1992 issue we published an article entitled "First IETF Internet Audiocast." The article described one of the first public demonstrations of a new class of Internet services that relies on IP multicast technologies. The authors, Steve Casner and Steve Deering wrote: "This event was a pilot experiment that we hope will be expanded at future IETF meetings to reach more destinations and to include video, images and shared whiteboards along with audio. This is a step toward a more distributed IETF, a goal Dave Farber and Jack Haverty challenged the IETF community to pursue during a discussion on the IETF mailing list last fall." In the past four years, the Internet community has come a long way in developing teleconferencing tools, and multicasting facilities are now a standard part of every IETF meeting. Our first article, by Mark Handley and Jon Crowcroft of University College London, describes the underlying architecture for Internet Multimedia Conferencing.

Access to the campus network, and thus the Internet, is becoming a requirement for students and staff at universities everywhere. In the past, such access could be provided by dialup modem pools which connected directly to mainframes or terminal servers. Today, the situation is much more complex and new strategies for campus networks are being developed. David Wasley of the University of California at Berkeley is responsible for the development of the Berkeley campus data network and associated services. In an article entitled "Building the Urban Area Network," he describes how the current campus network might evolve over the next few years.

The Telecommunications Reform Act of 1996 was signed into law by President Clinton on February 8th. The act makes a number of important changes to existing telecommunications law and it is expected to have a dramatic effect on the structure of broadcasting, cable television, and the telecommunications industries. Included in this new legislation is the *Communications Decency Act* (CDA). The CDA, is intended to protect children from having access to certain materials online, and makes it a Federal crime to provide such material electronically to anyone under the age of 18. The act has sparked a lot of debate amongst Internet users and providers, and raises many questions about the appropriate use of technology to limit or control content distribution. Included in this month's issue is an opinion piece about the CDA by Internet veteran John S. Quarterman.

Finally, if you receive this issue at NetWorld+Interop 96 Tokyo, we encourage you to take advantage of the special conference discount rate and sign up for a subscription to *ConneXions,* the official technical journal of NetWorld+Interop.

The Internet Multimedia Conferencing Architecture

by Mark Handley & Jon Crowcroft, University College London

Introduction

This article provides an overview of multimedia conferencing on the Internet. The protocols mentioned are all specified elsewhere as Internet-Drafts or RFCs. Each RFC gives details of the protocol itself, how it works and what it does. This article attempts to provide the reader with an overview of how the components fit together and some of the assumptions made.

The term "conferencing" is used in two different ways: firstly, to refer to bulletin boards and mailing list style *asynchronous* exchanges of messages between multiple users; secondly, to refer to *synchronous* or so-called "real-time" conferencing, including audio, video, shared whiteboards and other applications. This article is about the architecture for this latter application, in the Internet. There are other infrastructures for conferencing in the world: POTS (Plain Old Telephone System) networks often provide voice conferencing and phone-bridges, while the ISDN provides H.320 [1] for small, strictly organised video-telephony conferencing.

The architecture that has evolved in the Internet is far more general as well as being scalable to very large groups, and permits the open introduction of new media and new applications as they are devised.

There are a number of components to this architecture, and the rest of this article describes these as follows:

- Group Communication support: Multicast
- Data delivery models: reliability and performance assurance
- Transport Protocols: synchronisation and playout, feedback
- Conference Setup
- Security

The protocol stacks for Internet multimedia conferencing are shown in Figure 1. Most of the protocols are not deeply layered unlike many protocol stacks, but rather are used alongside each other to produce a complete conference.

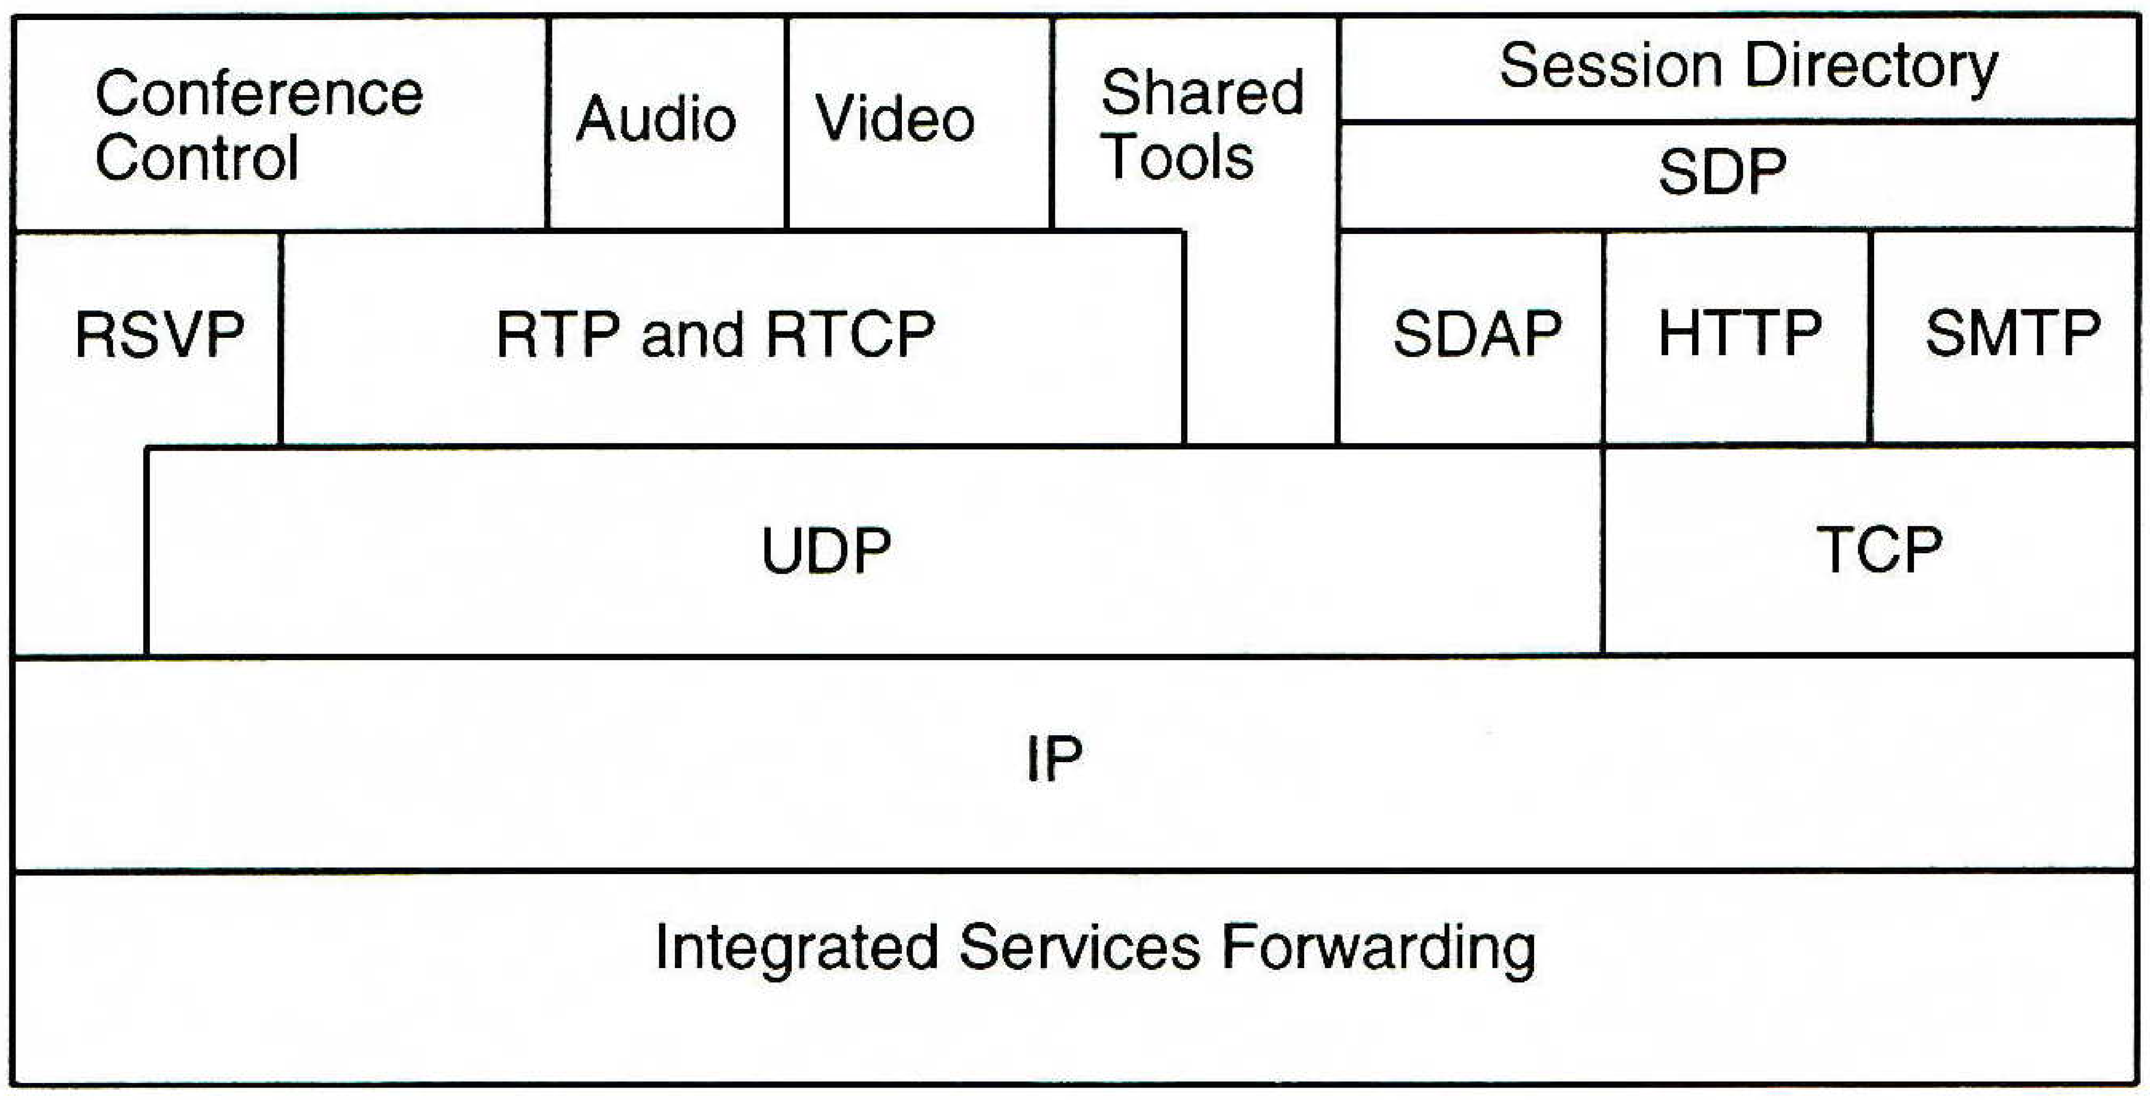

Figure 1: Internet Multimedia Conferencing Protocol Stacks

Multicast traffic distribution

IP multicast enables efficient many-to-many datagram distribution. It is one of the basic building blocks of the Internet multimedia conferencing architecture. For most conferencing purposes, unicast is viewed as being a special case of multicast routing.

Multicast Service Model

The IP multicast service model is as follows:

- Senders send datagrams to a multicast group address.
- Receivers express an interest in (join) certain multicast groups.
- Multicast routers conspire to deliver multicast group addressed datagrams from the senders to the receivers.

The important factor here is that senders do not have to know who the receivers are in order to be able to send to them. In fact, in most situations, no single point in the network needs to know who all the receivers are, and it is this that makes IP multicast scalable to very large groups. In addition, receivers do not need to know who the senders are in order to be able to receive traffic from them, and this solves many conference setup and resource location problems without needing explicit machinery.

There are many multicast routing protocols [2,3,4,5], but all of them satisfy the above service model. They differ in their mechanisms and in how they scale with number of senders and groups.

On a LAN, group membership is expressed by IGMP [6]. IGMP version 3 allows receivers to express an interest in only receiving some of the senders to a particular multicast group. Earlier versions of IGMP only allow a receiver to request to receive all the sources sending to a multicast group.

Address allocation

How does an application choose a multicast address to use? In the absence of any other information, we can bootstrap a multicast application by using *well-known* multicast addresses. Routing (unicast and multicast) and group membership protocols [6] can do just that. However, this is not the best way to manage applications of which there is more than one instance at any one time.

For these, we need a mechanism for allocating group addresses dynamically, and a directory service which can hold these allocations together with some key (session information for example—see later), so that users can look up the address associated with the application. The address allocation and directory functions should be distributed to scale well.

Address allocation schemes should avoid clashes, hence some kind of hash function suggests itself. Furthermore, both the address allocation system and the directory service can take advantage of the baseline multicast mechanism by advertising sessions through multicast messages on a *well-known* address, and using this to inform other directory servers to remove clashes and inform applications of the allocation.

Internet Service Models

Traditionally the Internet has provided best-effort delivery of datagram traffic from senders to receivers. No guarantees are made regarding when or if a datagram will be delivered to a receiver, however datagrams are normally only dropped when a router exceeds a queue size limit due to congestion. The best-effort Internet service model does not assume FIFO queuing, although many routers have implemented this.

With best-effort service, if a link is not congested, queues will not build at routers, datagrams will not be discarded in routers, and delays will consist of serialisation delays at each hop plus propagation delays. With sufficiently fast link speeds, serialisation delays are insignificant compared to propagation delays.

continued on next page

Internet Multimedia Conferencing *(continued)*

If a link is congested, with best-effort service queuing delays will start to influence end-to-end delays, and packets will start to be lost as queue size limits are exceeded.

Non-best effort service

Real-time Internet traffic is defined as datagrams that are delay sensitive. It could be argued that all datagrams are delay sensitive to some extent, but for these purposes we refer only to datagrams where exceeding an end-to-end delay bound of a few hundred milliseconds renders the datagrams useless for the purpose they were intended. For the purposes of this definition, TCP traffic is normally not considered to be real-time traffic, although there may be exceptions to this rule.

On congested links, best-effort service queuing delays will adversely affect real-time traffic. This does not mean that best-effort service cannot support real-time traffic—merely that congested best-effort links seriously degrade the service provided. For such congested links, a better-than-best-effort service is desirable.

Flows

To achieve this, the service model of the routers can be modified. At a minimum, FIFO queuing can be replaced by packet forwarding strategies that discriminate different "flows" of traffic. The idea of a flow is very general. A flow might consist of "all marketing site Web traffic," or "all fileserver traffic to and from teller machines" or "all traffic from the CEO's laptop wherever it is." On the other hand, a flow might consist of a particular sequence of packets from an application in a particular machine to a peer application in another particular machine between specific times of a specific day.

Flows are typically identifiable in the Internet by the tuple: {source machine, destination machine, source port, destination port, protocol} any of which could be "ANY" (wildcarded).

In the multicast case, the destination is the group, and can be used to provide efficient aggregation.

Flow identification is called *classification* and a class (which can contain one or more flows) has an associated service model applied. This can default to best effort.

Through network management, we can imagine establishing classes of long lived flows—enterprise networks ("Intranets") often enforce traffic policies that distinguish priorities which can be used to discriminate in favor of more important traffic in the event of overload (though in an under-loaded network, the effect of such policies will be invisible, and may incur no load/work in routers).

The router service model to provide such classes with different treatment can be as simple as a priority queuing system, or it can be more elaborate.

Although best-effort services can support real-time traffic, classifying real-time traffic separately from non-real-time traffic and giving real-time traffic priority treatment ensures that real-time traffic sees minimum delays. Non-real-time TCP traffic tends to be elastic in its bandwidth requirements, and will then tend to fill any remaining bandwidth.

We could imagine a future Internet with sufficient capacity to carry all of the world's telephony traffic. Since this is a relatively modest capacity requirement, it might be simpler to establish "POTS" as a static class which is given some fraction of the capacity overall, and then no individual call need be given an allocation (i.e., we would no longer need the call setup/tear down that was needed in the legacy POTS which was only present due to under-provisioning of trunks, and to allow the trunk exchanges the option of call blocking). The vision is of a network that is engineered with capacity for all of the average load sources to send all the time.

Reservations

For flows that may take a significant fraction of the network (i.e., are "special"), we need a more dynamic way of establishing these classifications. In the short term, this applies to any multimedia calls since the Internet is largely under-provisioned at the time of writing.

The *Resource ReserVation Protocol* (RSVP) is being standardised for just this purpose. It provides flow identification and classification. Hosts and applications are modified to speak the RSVP client language, and routers speak RSVP. [15]

Since most traffic requiring reservations is delivered to groups (e.g., TV), it is natural for the receiver to make the request for a reservation for a flow. This has the added advantage that different receivers can make heterogeneous requests for capacity from the same source. Thus RSVP can accommodate monochrome, color and HDTV receivers from a single source. Again the routers conspire to deliver the right flows to the right locations. RSVP accommodates the wildcarding noted above.

Admission control

If a network is provisioned such that it has excess capacity for all the real-time flows using it, a simple priority classification ensures that real-time traffic is minimally delayed. However, if a network is insufficiently provisioned for the traffic in a real-time traffic class, then real-time traffic will be queued, and delays and packet loss will result. Thus in an under-provisioned network, either all real-time flows will suffer, or some of them must be given priority.

RSVP provides a mechanism by which an admission control request can be made, and if sufficient capacity remains in the requested traffic class, then a reservation for that capacity can be put in place.

If insufficient capacity remains, the admission request will be refused, but the traffic will still be forwarded with the default service for that traffic's traffic class. In many cases even an admission request that failed at one or more routers can still supply acceptable quality as it may have succeeded in installing a reservation in all the routers that were suffering congestion. This is because other reservations may not be fully utilising their reserved capacity.

Billing

If a reservation involves setting aside resources for a flow, this will tie up resources so that other reservations may not succeed, and depending on whether the flow fills the reservation, other traffic is prevented from using the network. Clearly some negative feedback is required in order to prevent pointless reservations from denying service to other users. This feedback is typically in the form of billing. For real-time non-best effort traffic that is not reserved, this negative feedback is provided in the form of loss due to congestion of a traffic class, and it is not clear that usage based billing is required.

continued on next page

Internet Multimedia Conferencing *(continued)*

Billing requires that the user making the reservation is properly authenticated so that the correct user can be charged. Billing for reservations introduces a level of complexity to the Internet that has not typically been experienced with non-reserved traffic, and requires network providers to have reciprocal usage-based billing arrangements for traffic carried between them. It also requires mechanisms whereby some fraction of the bill for a link reservation can be charged to each of the downstream multicast receivers.

IP over ATM and RSVP

When IP operates over ATM, we can envisage a problem—the RSVP signaling model is receiver based, and the IP distribution model is many-to-many.

IP multicast over ATM point to multipoint distribution has been addressed by the IP-ATM working group, and at least two solutions (MARS, and PIM-SM or CBT over point-multipoint UNI 3.1 VCs) are feasible right now. This area is one where there is a lot of activity, and liaison between the IETF and the ATM Forum is helping resolve some of the problems.

The signaling is more complex—one solution has been proposed by Berson and is illustrated in Figure 2.

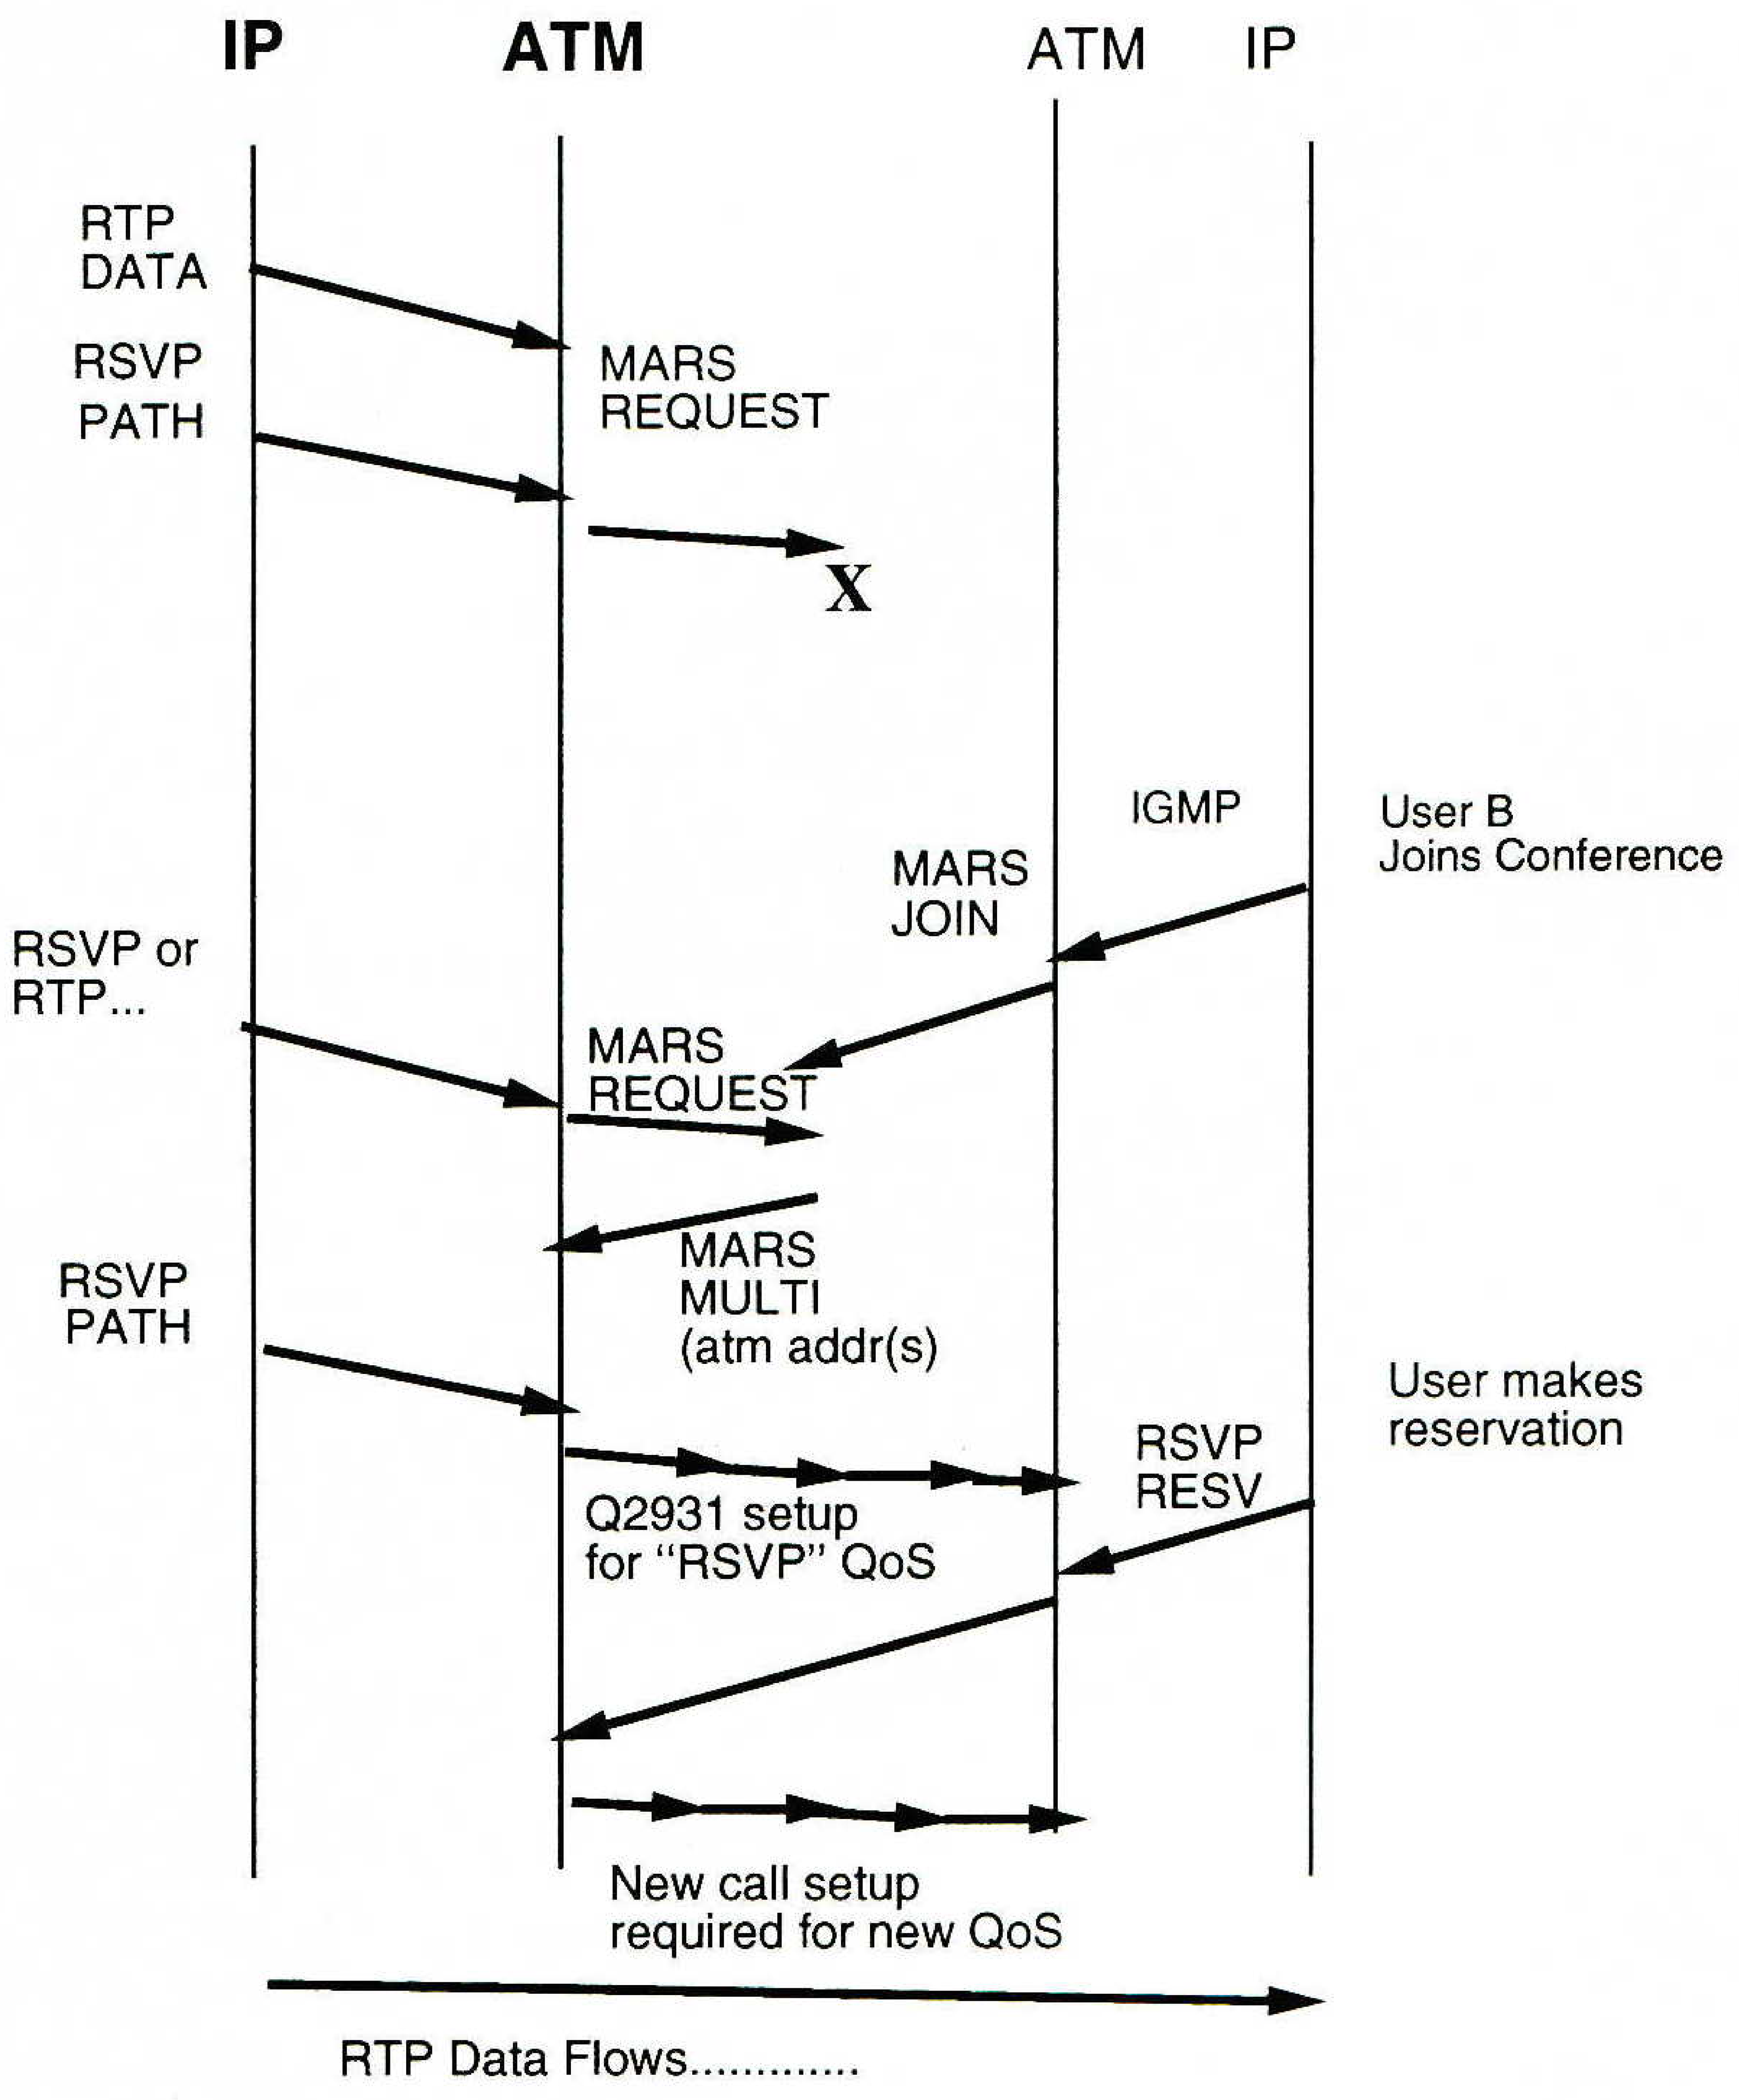

Figure 2: Berson's RSVP to Q2931 Interworking

Looking at this in more detail we can see that there is a lot of overhead. However, in Figure 3, we show how this is very much simplified when the ATM cloud is capable of leaf joins and QoS renegotiation.

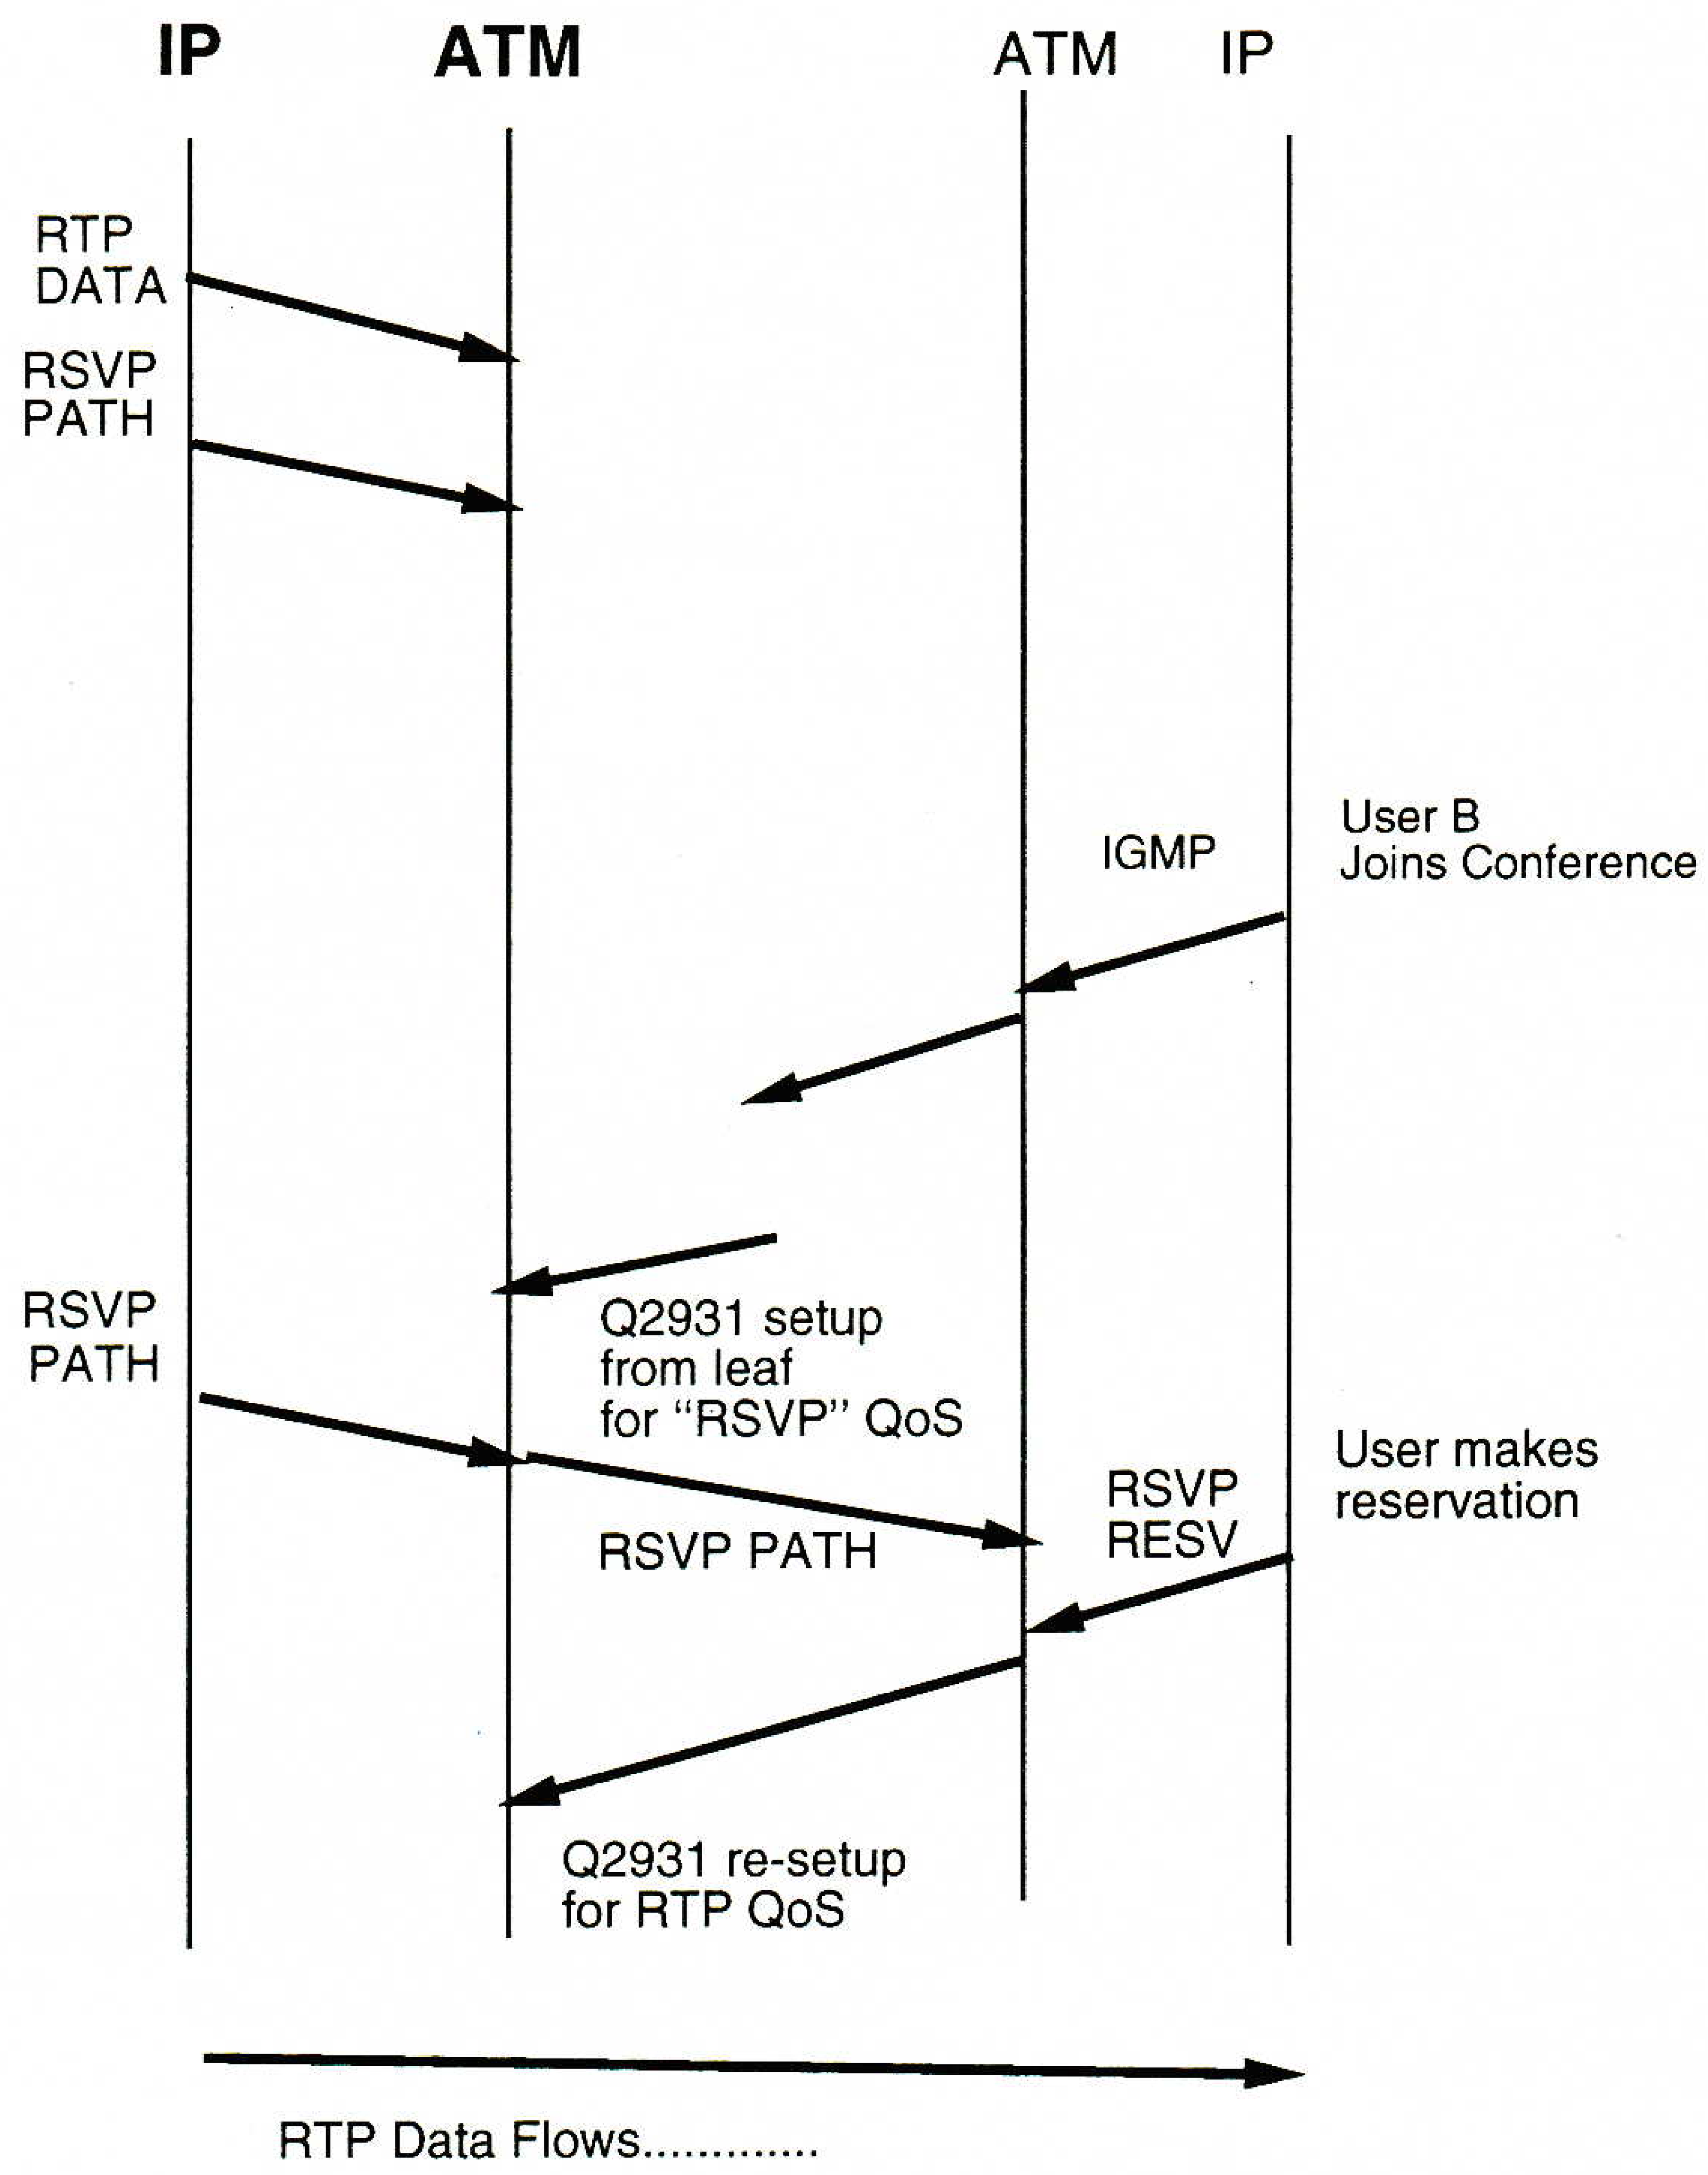

Figure 3: RSVP to Q2931 with Leaf Join and Renegotiation

Ongoing work in the ATM community is adding "leaf join" signaling to point-multipoint calls, and even extending the QoS model to allow renegotiation. Eventually, we might see an interworking model that is nearly transparent as shown in Figure 4. The goal is to include the path selection and QoS in both the IP and ATM worlds in one architecture.

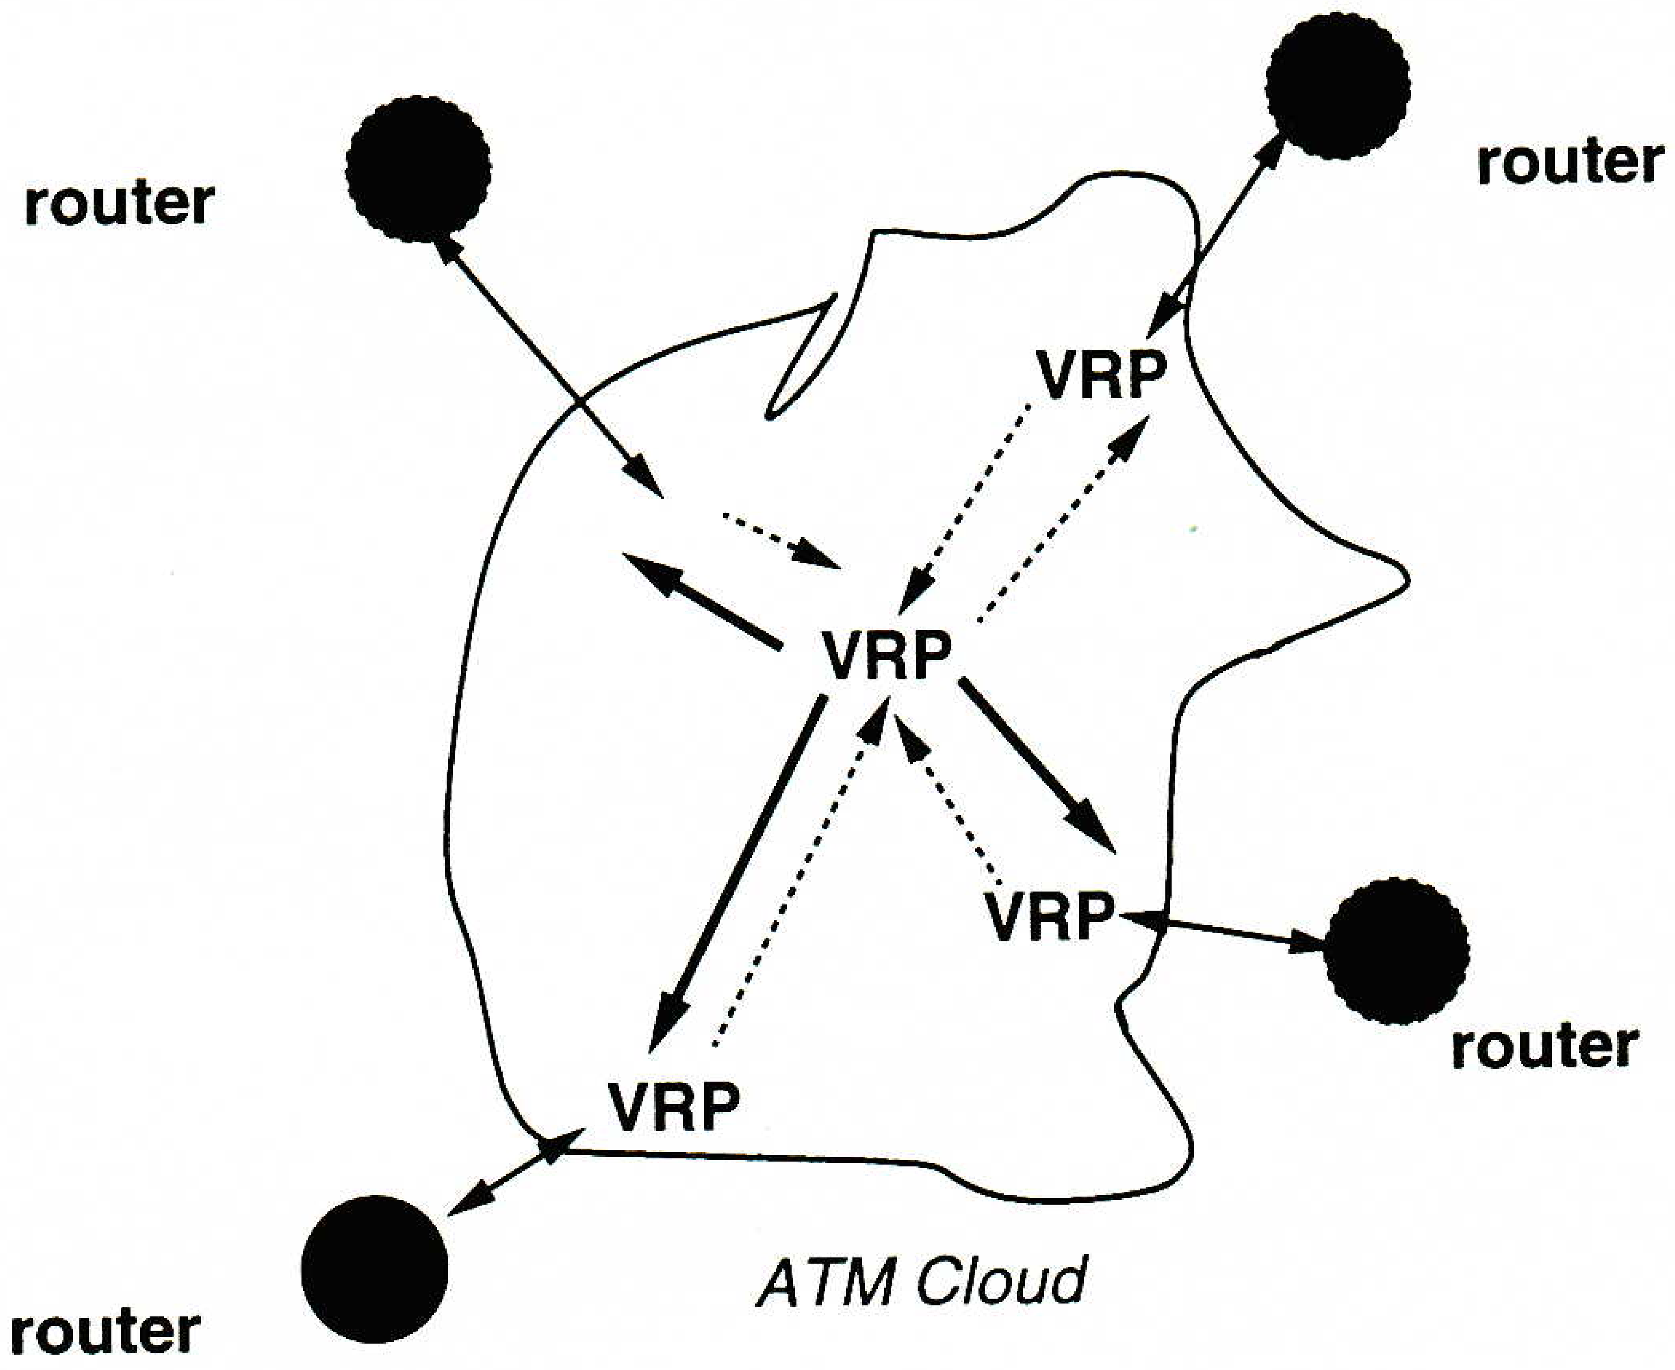

Figure 4: Transparent RSVP to Q2931 multicast Interworking

continued on next page

Internet Multimedia Conferencing *(continued)*

This is a simplification, but at least gives one a picture that the flexibility and low cost of the Internet approach may not have to vanish when operating over ATM! Here the use of *Virtual Rendezvous Points* (VRP) inside the ATM cloud to *attract* multicast sending and rendezvous receivers is shown. The scheme to select virtual RPs is a combination of the ATM Multicast technology, MARS, and the well known hash function in hierarchical PIM. Packets are then unicast from sources into the ATM cloud and up to the MCS/VRP and multicast out to exit points of the cloud, and on into the Internet.

Transport protocols

So-called *real-time delivery* of traffic requires little in the way of a transport protocol. In particular, real-time traffic that is sent over more than trivial distances is not re-transmittable.

Separate flows for each Media Stream

With packet multimedia data there is no need for the different media comprising a conference to be carried in the same packets. In fact it simplifies receivers if different media streams are carried in separate flows (i.e., separate transport ports and/or separate multicast groups). This also allows the different media to be given different quality of service. For example, under congestion, a router might preferentially drop video packets over audio packets. In addition, some sites may not wish to receive all the media flows. For example, a site with a slow access link may be able to participate in a conference using only audio and a whiteboard whereas other sites in the same conference may also send and receive video.

Receiver adaptation

Best-effort traffic is delayed by queues in routers between the sender and the receivers. Even reserved priority traffic may see small transient queues in routers, and so packets comprising a flow will be delayed for different times. Such delay variance is known as *jitter*.

Real-time applications such as audio and video need to be able to buffer real-time data at the receiver for sufficient time to remove the jitter added by the network and recover the original timing relationships between the media data. In order to know how long to buffer for, each packet must carry a timestamp which gives the time at the sender when the data was captured. Note that for audio and video data timing recovery, it is not necessary to know the absolute time that the data was captured at the sender, only the time relative to the other data packets.

Synchronisation

As audio and video flows will receive differing jitter and possibly differing quality of service, audio and video that were grabbed at the same time at the sender may not arrive at the receiver at the same time. At the receiver, each flow will need a playout buffer to remove network jitter. Inter-flow synchronisation can be performed by adapting these playout buffers so that samples/frames that originated at the same time are played out at the same time. This requires that the times that different flows from the same sender were captured are available at the receivers.

RTP

The transport protocol for real-time flows is RTP [7]. This provides a standard format packet header which gives media specific timestamp data, as well as payload format information and sequence numbering amongst other things. RTP is normally carried using UDP. It does not provide or require any connection setup, nor does it provide any enhanced reliability over UDP. For RTP to provide a useful media flow, there must be sufficient capacity in the relevant traffic class to accommodate the traffic. How this capacity is ensured is independent of RTP.

RTP media timestamps units are flow specific—they are in units that are appropriate to the media flow. For example, 8kHz sampled PCM encoded audio has a timestamp clock rate of 8kHz. This means that inter-flow synchronisation is not possible from the RTP timestamps alone.

Every original RTP source is identified by a source identifier, and this source ID is carried in every packet. RTP allows flows from several sources to be mixed in gateways to provide a single resulting flow. When this happens, each mixed packet contains the source IDs of all the contributing sources.

Each RTP flow is supplemented by *Real-Time Control Protocol* (RTCP) packets. There are a number of different RTCP packet types. RTCP packets provide the relationship between the realtime clock at a sender and the RTP media timestamps, and provide textual information to identify a sender in a conference from the source ID.

Conference membership and reception feedback

IP multicast allows sources to send to a multicast group without being a receiver of that group. However, for many conferencing purposes it is useful to know who is listening to the conference, and whether the media flows are reaching receivers properly. Accurately performing both these tasks restricts the scaling of the conference. IP multicast means that no-one knows the precise membership of a conference at a specific time, and this information cannot be discovered, since trying to do so would cause an implosion of messages, many of which would be lost. (This is not to say that we cannot know the bounds of a conference membership, a subset of whom might be present at any time—this can be done using encryption and restricted distribution of encryption keys, or which more later). Instead, RTCP provides approximate membership information through periodic multicast of session messages which, in addition to information about the recipient, also give information about the reception quality at that receiver. RTCP session messages are restricted in rate, so that as a conference grows, the rate of session messages remains constant, and each receiver reports less often. A member of the conference can never know exactly who is present at a particular time from RTCP reports, but does have a good approximation to the conference membership.

Reception quality information is primarily intended for debugging purposes, as debugging of IP multicast problems is a difficult task. However, it is possible to use reception quality information for rate adaptive senders, although it is not clear whether this information is sufficiently timely to be able to adapt fast enough to transient congestion. However, it is certainly sufficient for Van Jacobson Congestion Control style adaption to a "share" of the current capacity.

Conference setup

Conferences come in many shapes and sizes, but there are only really two models for conference control: light-weight sessions and tightly coupled conferencing. For both models, a rendezvous mechanism is needed. Note that the conference control model is orthogonal to issues of quality of service and network resource reservation.

Light-weight sessions

Light-weight sessions are multicast based multimedia conferences that lack explicit session membership and explicit conference control mechanisms. Typically a lightweight session consists of a number of many-to-many media streams supported using RTP and RTCP using IP multicast.

continued on next page

Internet Multimedia Conferencing *(continued)*

The rendezvous mechanism for light-weight sessions is a multicast based session directory. This distributes session descriptions [8] to all the potential session participants. These session descriptions provide an advertisement that the session will exist, and also provide sufficient information including multicast addresses, ports, media formats and session times so that a receiver of the session description can join the session. As dynamic multicast address allocation can be optimised by knowing which addresses are in use at which times, the session directory is an appropriate agent to perform multicast address allocation.

Tightly coupled conferences

Tightly coupled conferences may also be multicast based and use RTP and RTCP, but in addition they have an explicit conference membership mechanism and may have an explicit conference control mechanism that provides facilities such as floor control.

Such conferences may be initiated either by invitation (the "conference" calls the user), or by user initiation (the user calls the "conference"). In the latter case the rendezvous mechanism can be handled by the same session directory that handles light-weight sessions, with the addition of a description of the contact mechanism to be used to join the conference to the description of the session. In the former case, a call up mechanism is required which can be combined with the explicit conference membership mechanism.

No standard mechanism currently exists to perform either the conference membership mechanism or the "dial-up" mechanism in the Internet, and the many proprietary conferencing systems available all implement this in different ways. At the time of writing, it seems likely that a protocol based on the ITU's T.124 [9] recommendation will be derived for Internet usage.

Security

There is a temptation to believe that multicast is inherently less private than unicast communication since the traffic visits so many more places in the network. In fact, this is not the case except with broadcast and prune type multicast routing protocols [3]. However, IP multicast does make it simple for a host to anonymously join a multicast group and receive traffic destined to that group without the other senders' and receivers' knowledge. If the application requirement is to communicate between some set of users, then strict privacy can only be enforced in any case through adequate end-to-end encryption.

RTP specifies a standard way to encrypt RTP and RTCP packets using private key encryption schemes such as DES [10]. It also specifies a standard mechanism to manipulate plain text keys using MD5 [11] so that the resulting bit string can be used as a DES key. This allows simple out-of-band mechanisms such as Privacy-Enhanced Mail (PEM) to be used for encryption key exchange.

Authentication and Key Distribution

Key distribution is closely tied to authentication. Conference or session directory keys can be securely distributed using public-key cryptography on a one-to-one basis (by e-mail, a directory service, or by an explicit conference setup mechanism), but this is only as good as the certification mechanism used to certify that a key given by a user is the correct public key for that user. Such certification mechanisms [12] are not specific to conferencing, and no standard mechanisms are currently in use for conferencing purposes other than PEM [13].

Even without privacy requirements, strong authentication of a user is required if making a network reservation results in usage based billing.

Encrypted session Announcements

Session Directories can make encrypted session announcements using private key encryption, and carry the encryption keys to be used for each of the conference media streams in the session. Whilst this does not solve the key distribution problem, it does allow a single conference to be announced more than once to more than one key-group, where each group holds a different session directory key, so that the two groups can be brought together into a single conference without having to know each other's keys.

Session Directories and Invitiation

Recent work in the IETF MMUSIC group has produced specifications for the Session Directory, and now for a Session Invitation Protocol. Until the advent of these, the Mbone has been somewhat akin to Citizen Band Radio: groups of users somewhat anarchically tuning in and listening or sending on dynamically randomly allocated addresses.

To provide some level of coordination, the session directory was designed so that users can coordinate the allocation of addresses to named sessions, and disseminate information about a session (media types, start and end times, related information on the World-Wide Web, and so on can all be sent out in a session advertisement). Session advertisements are sent out in UDP multicast packets to a well known multicast address, by daemons periodically, and cached when users run listeners. Users can create sessions through a GUI, and they can even browse sessions from a Web browser.

The session attribute syntax is quite complex, and this is not the place to go into it, but it was observed, that a relatively simple addition of user location would enable the same information to be used to invite users to join a session. This has led to work on a *Session Invitation Protocol* based on a simple model somewhat akin to a combination of a telephone, answerphone and call redirector—this is work in progress, so it is too early to say more about this here, but it should lead to convergence amongst the many developers of Internet Phone software, in the area of user-to-user signaling.

Applications other than audio and video

There are a lot of other applications in use alongside video and audio in multimedia conferencing systems. Three notable ones are *Imm, Wb* [14] and *Nt.*

Imm is used to send out large images (weather satellite ones in particular). *Wb* is used for a shared whiteboard, and is familiar to most Mbone users. *Nt* is a new tool from UCL that allows multiple people to work on editing a text document at the same time. This is aimed at distributed, replicated minute keepers or document drafting type meetings, whereas *Wb* is more useful for brainstorming.

Examples

Figure 5 on the next page shows the time sequence involved in setting up a light-weight session between two sites. In this case, site A creates a session advertisement, and some time later starts sending a media stream even though there may be no receiver at that time.

Some time later, site B joins the session, and starts to receive the traffic. At the earliest opportunity site B also makes an RSVP reservation to ensure the flow quality is satisfactory.

continued on next page

Internet Multimedia Conferencing *(continued)*

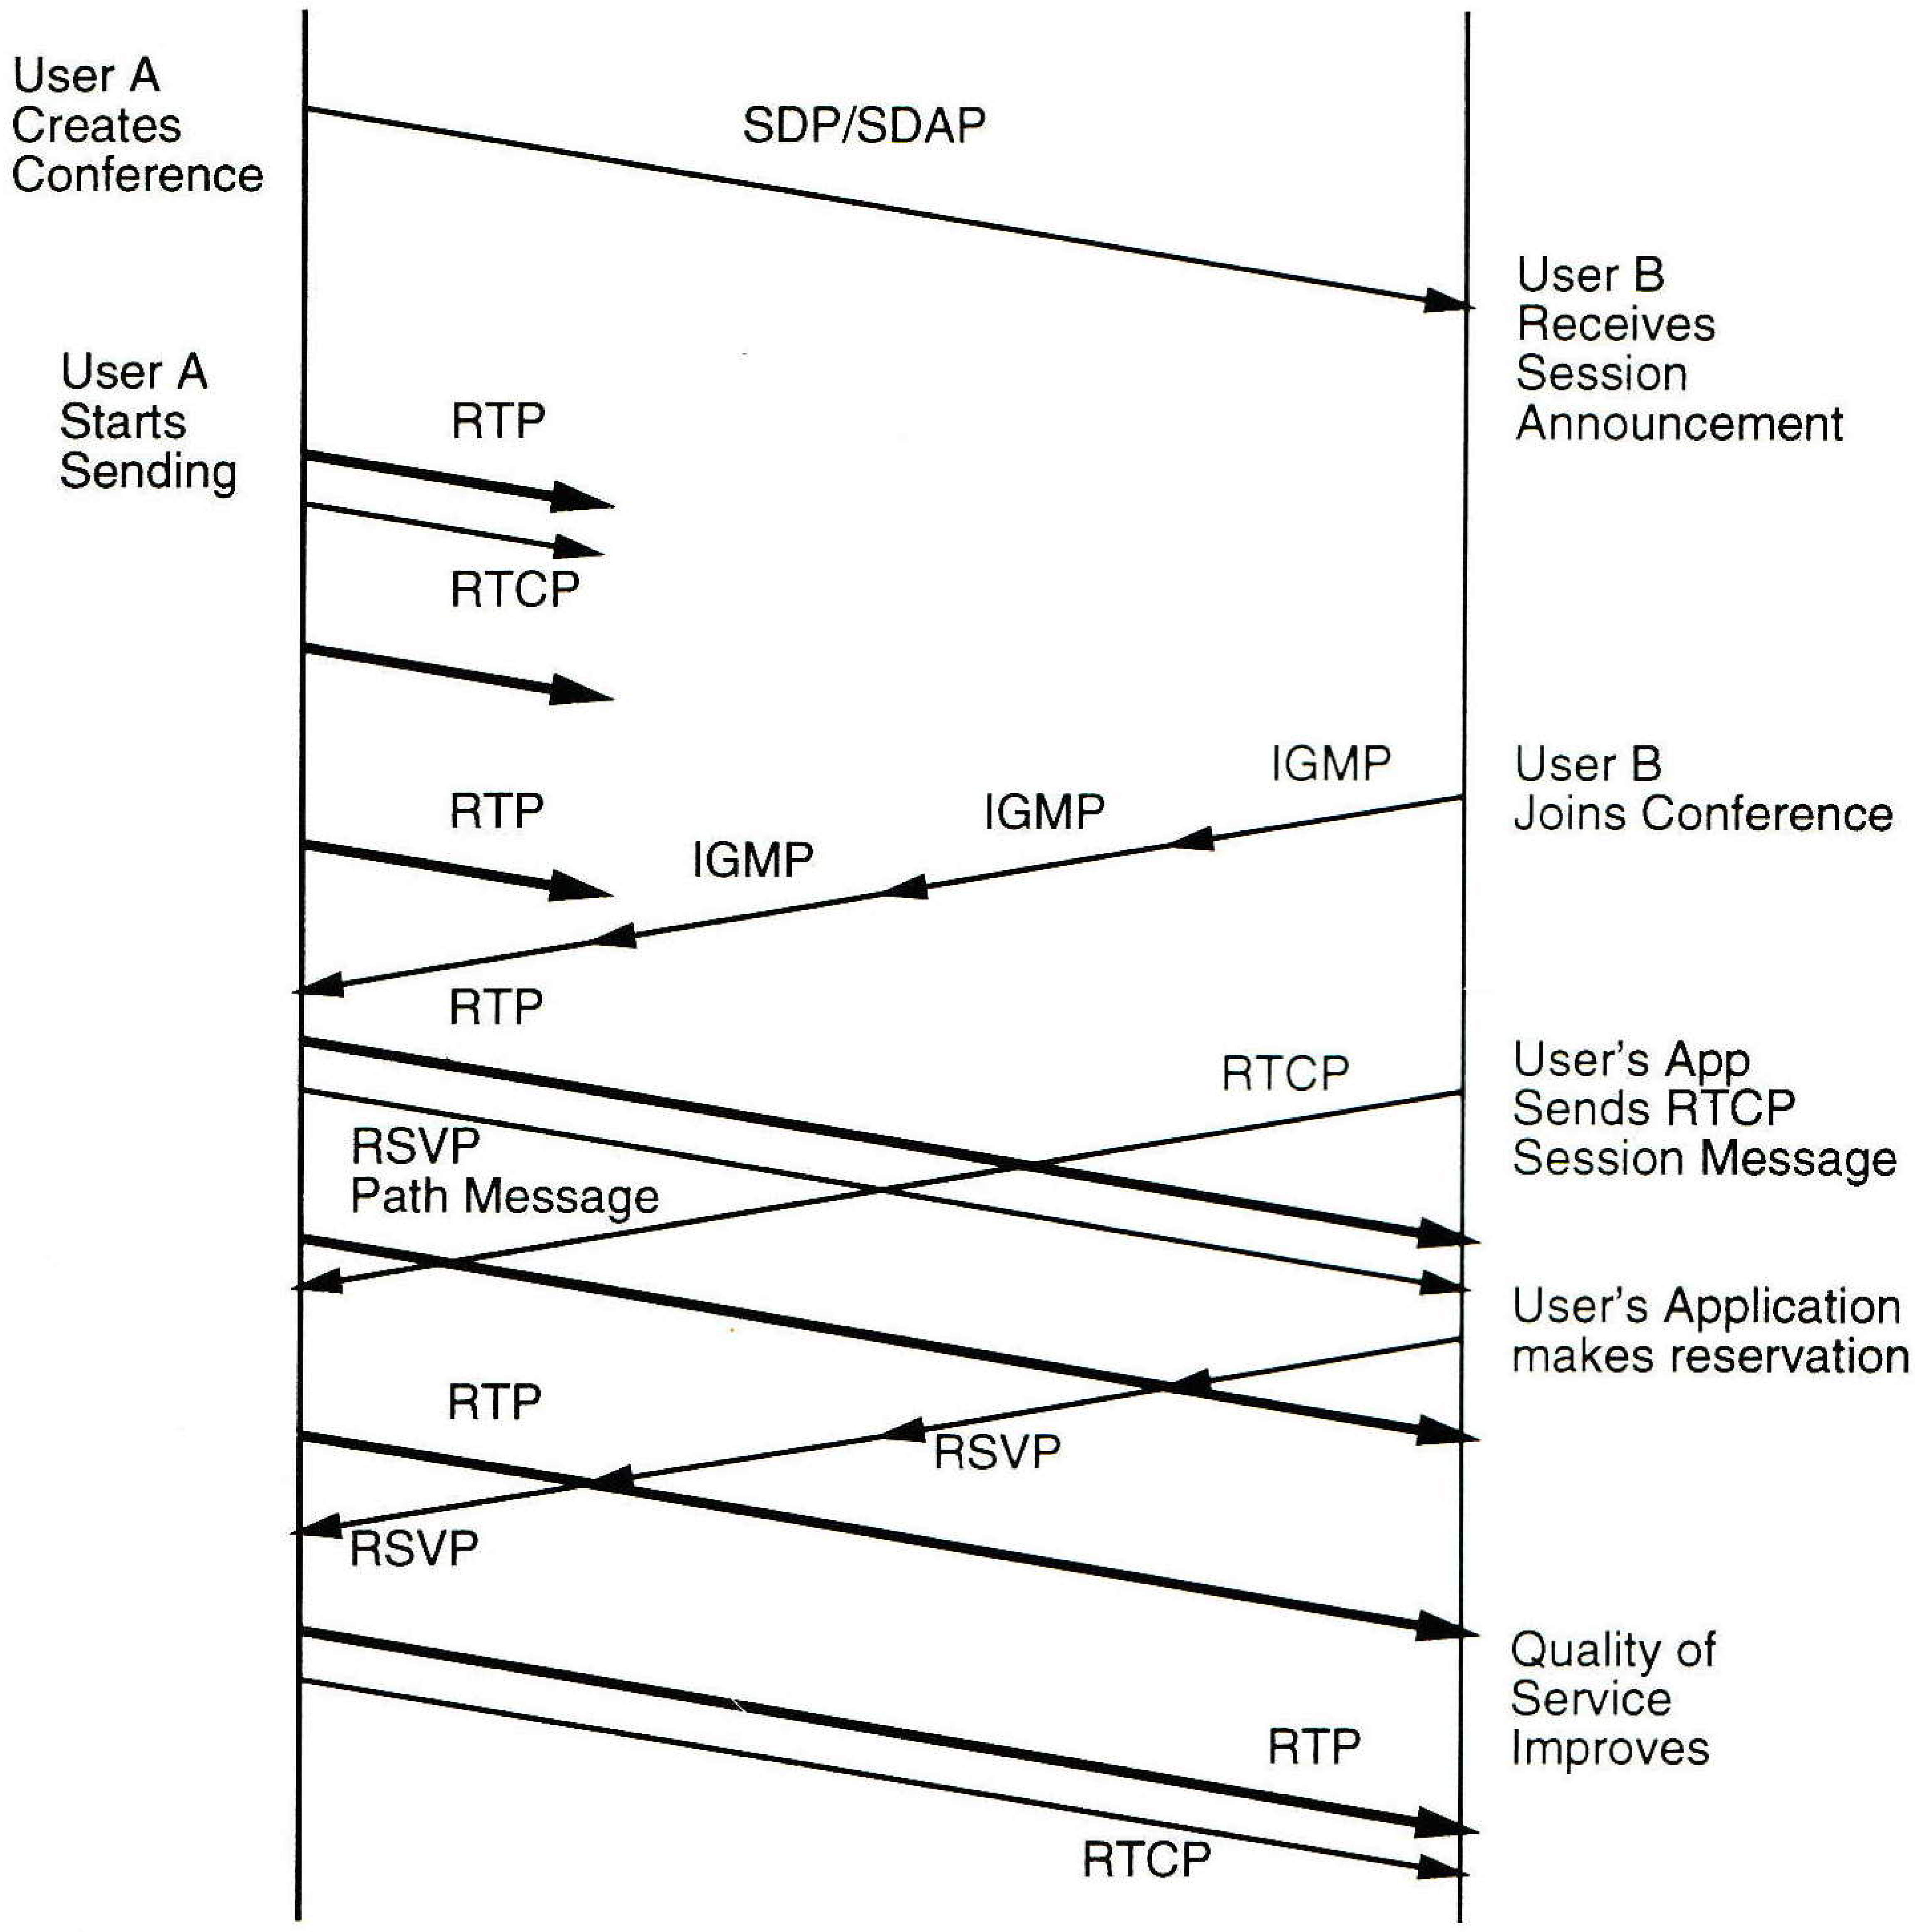

Figure 5: Joining a light-weight multimedia session

Summary/Conclusions

This article is an attempt to gather together in one place the set of assumptions behind the design of the Internet Multimedia Conferencing Architecture, and the ideas and services that are provided to support it.

This area is one of ongoing work, but we hope that we have clarified the basic goals of the work, and shown how there actually is a big picture within which it fits, which some observers may not have hitherto believed!

Acknowledgements

Acknowledgements are due to the End-to-End Research Group, the Int-serv, RSVP, MMUSIC and AVT working groups of the IETF, and discussion with colleagues at UCL. The earliest clear exposition of the ideas here can be found at:

`http://www-mice.cs.ucl.ac.uk/mice-old/van/`

...and was presented at ACM SIGCOMM 1994 in London by Van Jacobson.

Authors'address

Department of Computer Science
University College London
Gower Street
London WC1E 6BT
United Kingdom
Fax: +44 171 387 1397
Web: `http://www.cs.ucl.ac.uk/index.html`

References

[1] ITU Recommendation H.320

[2] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C-G. Liu, L. Wei, "An Architecture for Wide Area Multicast Routing," ACM SIGCOMM 1994, London October 1994, ACM *Computer Communications Review,* Volume 24, No. 4, pp 126–135.

[3] S. Deering, C. Partridge, D. Waitzman, "Distance Vector Multicast Routing Protocol," RFC 1075, November 1988.

[4] A. Ballardie, P. Francis, J. Crowcroft, "An Architecture for Scalable Inter-Domain Multicast Routing," Proceedings of ACM SIGCOMM 1993, pp 85–95.

[5] J. Moy, "Multicast Extensions to OSPF," RFC 1584, March 1994.

[6] Steve Deering, "Multicast Routing in Internetworks and Extended LANs," ACM SIGCOMM 88, August 1988, pp 55–64. See also "Host Extensions for IP Multicasting," RFC 1112, August 1989.

[7] H. Schulzrinne, S. Casner, R. Frederick and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications," Internet-Draft, Work In Progress, Late 1995.

[8] M. Handley, V. Jacobson, "SDP: Session Description Protocol," Internet-Draft, Work in Progress, November 1995.

[9] ITU Recommendation T.124, "Generic Conference Control."

[10] National Institute of Standards and Technology (NIST), "FIPS Publication 46-1: Data Encryption Standard," January 22, 1988.

[11] Rivest, R., "The MD5 Message-Digest Algorithm," RFC 1321, April 1992.

[12] CCITT (Consultative Committee on International Telegraphy and Telephony), "Recommendation X.509: The Directory—Authentication Framework," 1988.

[13] J. Linn, "Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures," RFC 1421, February 1993.

[14] S. Floyd, V. Jacobson, S. McCanne, C-G. Liu, L. Zhang, "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing," Proceedings of ACM SIGCOMM 1995.

[15] Braden, R., Zhang, L., "RSVP: A Resource ReserVation Protocol," *ConneXions,* Volume 8, No. 8, August 1994.

[16] Flückiger, F., "Back to Basics: Networking requirements of audio and motion video," *ConneXions,* Volume 10, No. 1, January 1996.

JON CROWCROFT is a senior lecturer in the Department of Computer Science, University College London, where he is responsible for a number of European and US funded research projects in Multi-media Communications. He has been working in these areas for over 14 years. He graduated in Physics from Trinity College, Cambridge University in 1979, and gained his MSc in Computing in 1981, and PhD in 1993. He is a member of the ACM, the BCS and the IEE. He is also on the editorial teams for the *Transactions on Networks* and the *Journal of Internetworking.* E-mail: `J.Crowcroft@cs.ucl.ac.uk`

MARK HANDLEY graduated from UCL with a 1st class honours degree in Computer Science with Electrical Engineering in 1988. As a PhD student at UCL, he studied novel neural network models and their visualisation. Since 1992, he has been a Research Fellow, working on the RACE CAR project and on MICE, now managing the UCL part of MICE. His current research interests include Multimedia Systems, especially audio and video encoding and compression, Distributed and Heterogeneous Systems, and HCI and graphics. E-mail: `M.Handley@cs.ucl.ac.uk`

Building the Urban Area Network

How do we ensure the infrastructure will exist when we need it?

by David L. Wasley, The University of California

Introduction

By the end of this decade most major universities will have a wide variety of on-line multimedia learning resources, digital libraries will be well underway, and use of the network will be a routine aspect of the educational process. On campus this presents no unusual problem, and the campus network service can be extended to university residence halls as well. The problem for which we must find a solution is to ensure that we will have in place by then a wide area communications infrastructure supporting affordable high capacity data services that will allow delivery of these essential resources to homes and other off campus locations. I believe the solution can be seen as an evolution from where we are now, through partnerships with new commercial service providers, towards the goal of being able to take full advantage of new developments in communications technology and infrastructure.

Where we are now

The Berkeley campus community includes around 44,000 faculty, staff, students and researchers. The campus network now includes over 20,000 nodes and is still growing at the rate of around 4,000 per year. Recent plans include support for 100 megabit/second LAN technologies and experimental installation of ATM in support of "video-on-demand" servers.

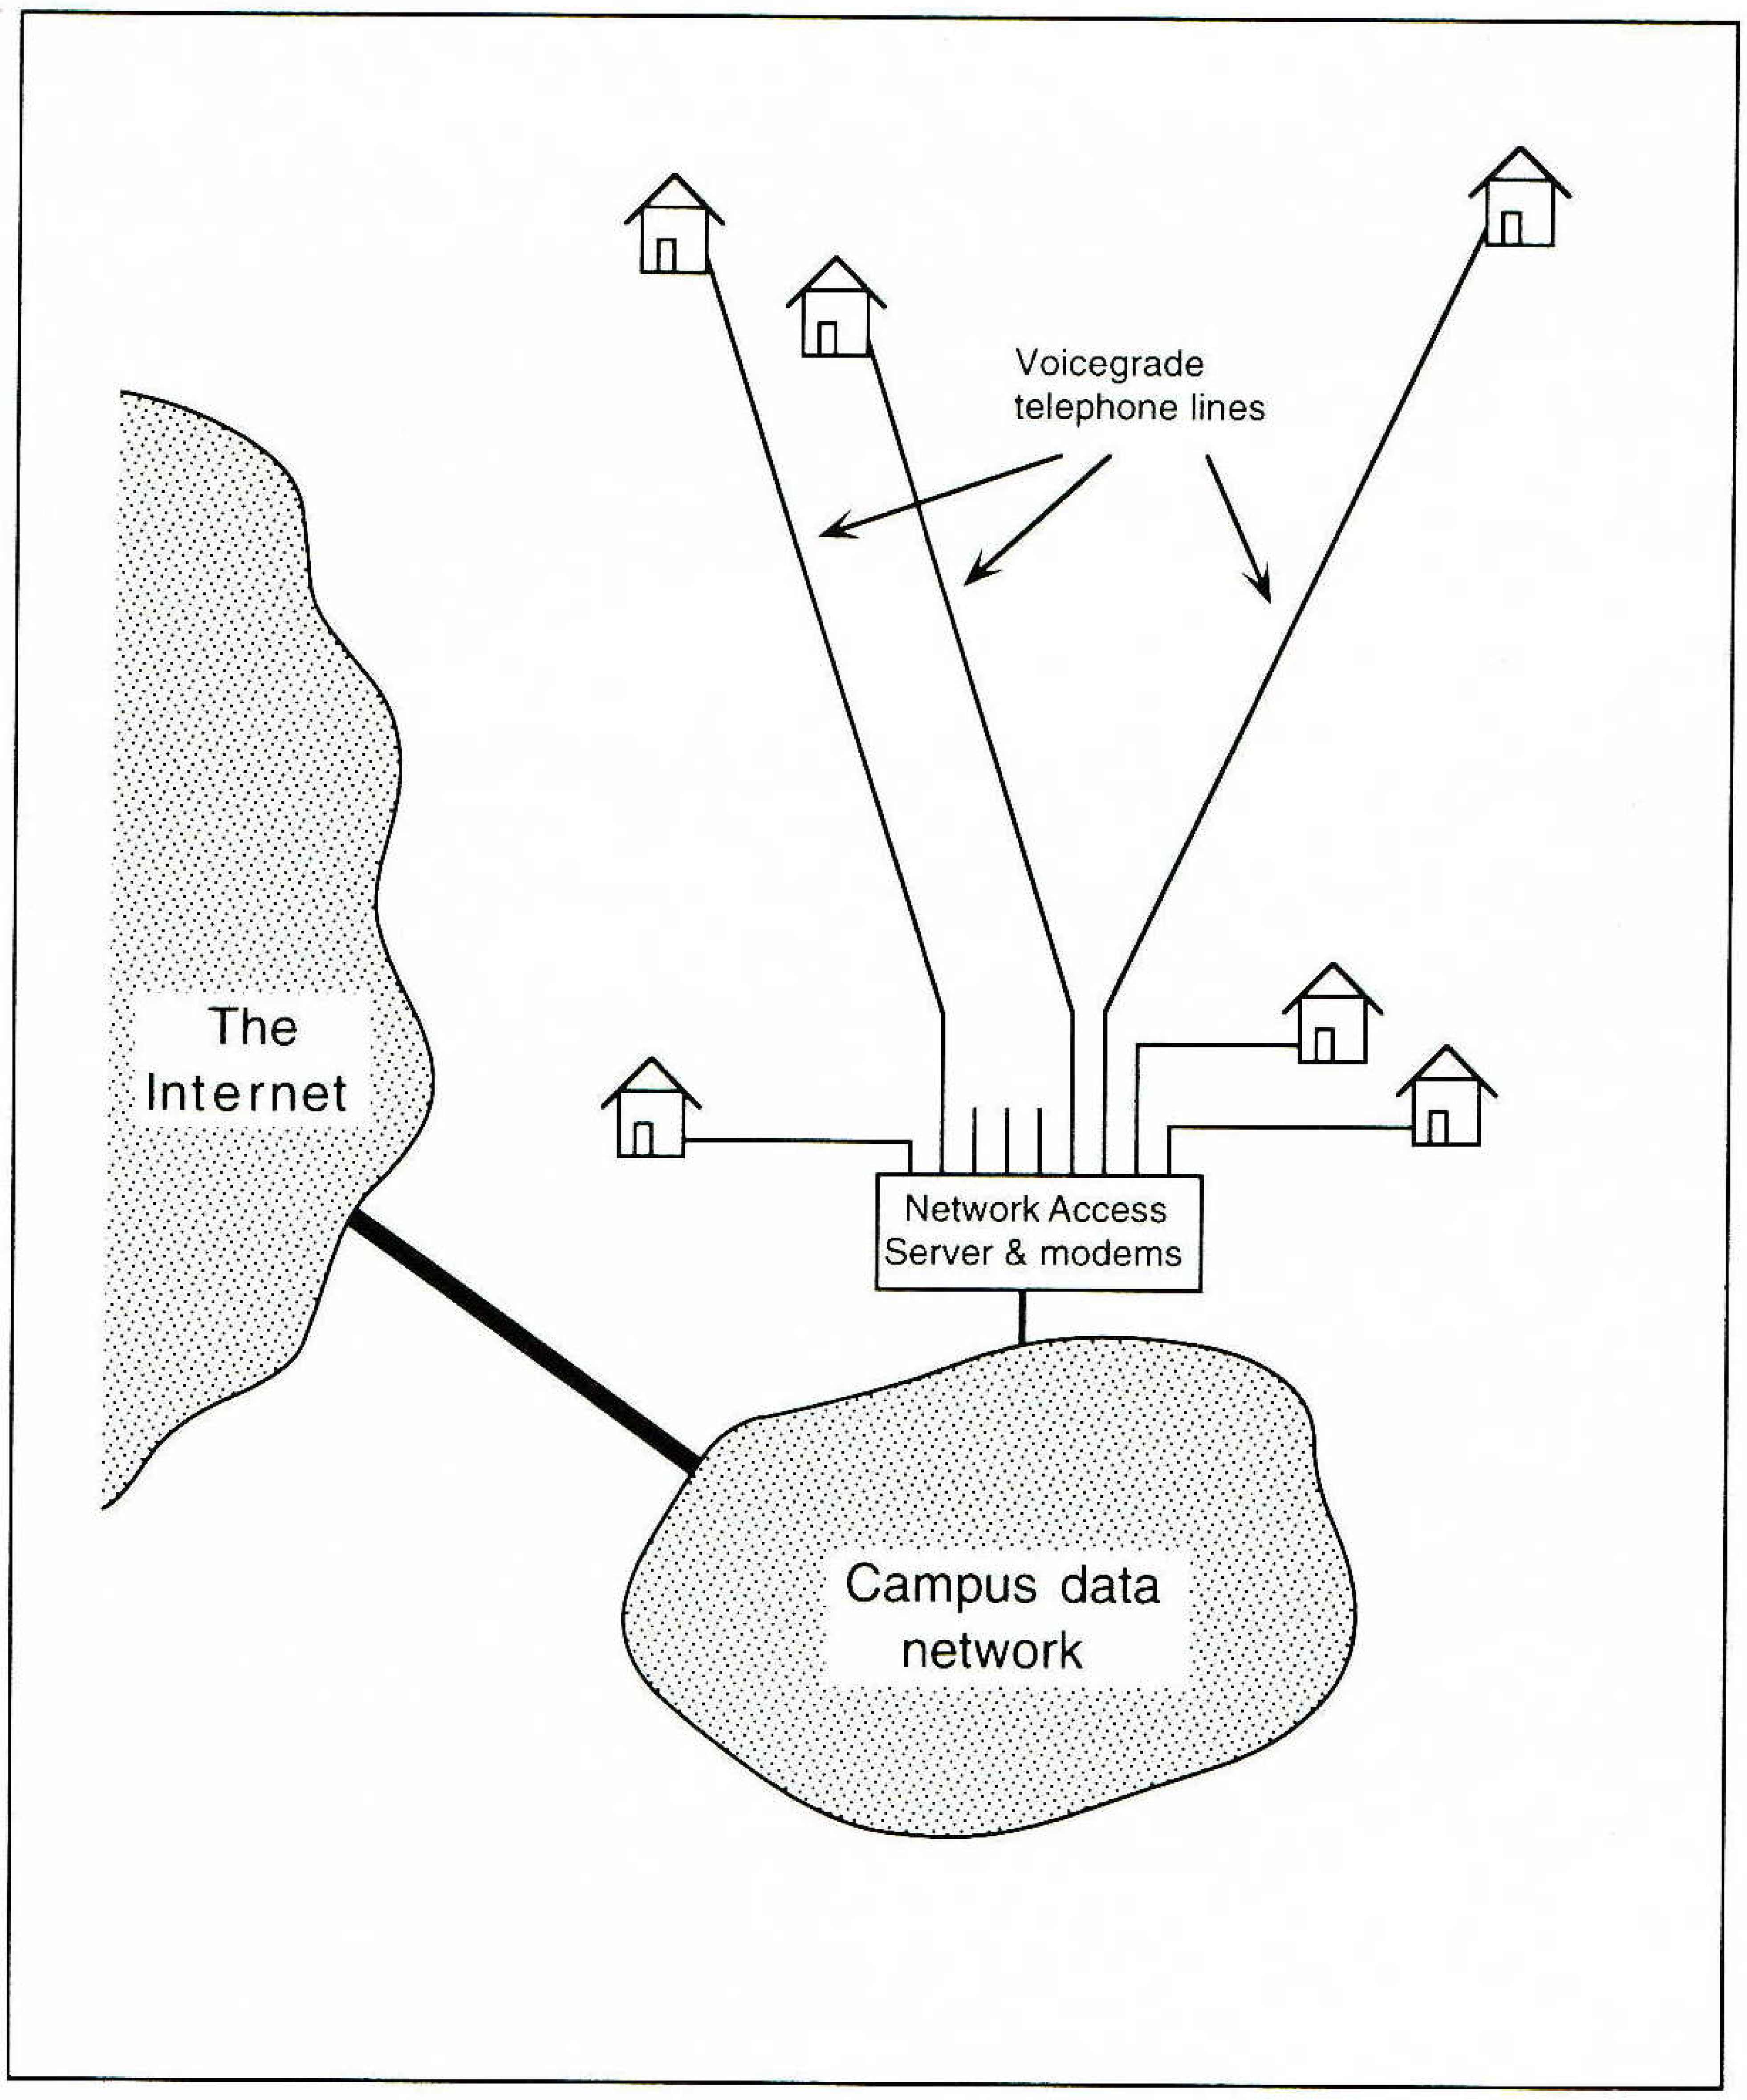

Figure 1: Traditional Method of Connecting to the Campus Network

The Data Communication & Network Services department at Berkeley now supports over 600 modems as part of its campus network services. (See Figure 1) Typically these modems are fully utilized from well before noon until after midnight. If we were to try to scale our modem service to meet projected demand, Berkeley would have many thousands of modems in place by the year 2000. We believe this would be a logistical and management nightmare so what alternative should we foster over the next half decade?

The use of network communications has become integral to our work, our research, and our teaching and learning activities. The need to maintain access to the network from a variety of off-campus locations has put increasing pressure on the Telecom and Datacom service units to provide more and more "dial-in" access points. Most campus members use electronic mail, gopher, and other information resources. Those that use a protocol based connection (PPP) also make heavy use of World-Wide Web (WWW) servers and campus-based file servers. Some have file servers at home and would like to use those resources while on campus. A few more technically oriented people have networks of computers in their homes from which they wish to communicate often with the campus network.

Historically most universities first installed dial-in modems as a "free" service, primarily in support of information systems staff and the occasional professor who needed to work at home. Today this service is used by everyone and demand is increasing very rapidly! The primary problem of expanding such a free or "library model" service is that of identifying a funding stream that scales with increasing demand. (Note: Libraries also are looking at this problem and many are considering charging for some services.)

Many campuses are converting to a billed modem service with a modest monthly fee. This incurs additional costs, of course, for account setup, more complex access control systems, billing data collection and processing, and for follow-up on delinquent subscribers.

- Is it reasonable for universities to develop this type of service?
- What alternatives might there be?

Most campus network systems are based on Internet technology (TCP/IP). All large campuses are connected to the world-wide Internet today. In the last year or so, dozens of "dial up" *Internet Service Providers* (ISPs) have set up business. Many are nation-wide in scope. A few ISPs even will set up modem equipment on your campus and operate it for your community. Most recently Sprint, MCI, and AT&T have joined the growing list of dial-up ISPs.

ISPs have a funding stream designed to allow them to support and expand their capacity (or they won't survive very long!). Therefore, why not "out source" access to the campus network by means of this new class of service providers? (See Figure 2) There are pros and cons to this strategy, but we believe that there are very good reasons to adopt this model as quickly as possible.

Why not use commercial ISPs?

A major concern on most campuses is total cost to the end-user. If access to the network is integral to the educational program, students will require it. If there is a cost involved, we must be concerned about equity issues and the overall cost of education to the student.

continued on next page

Building the Urban Area Network *(continued)*

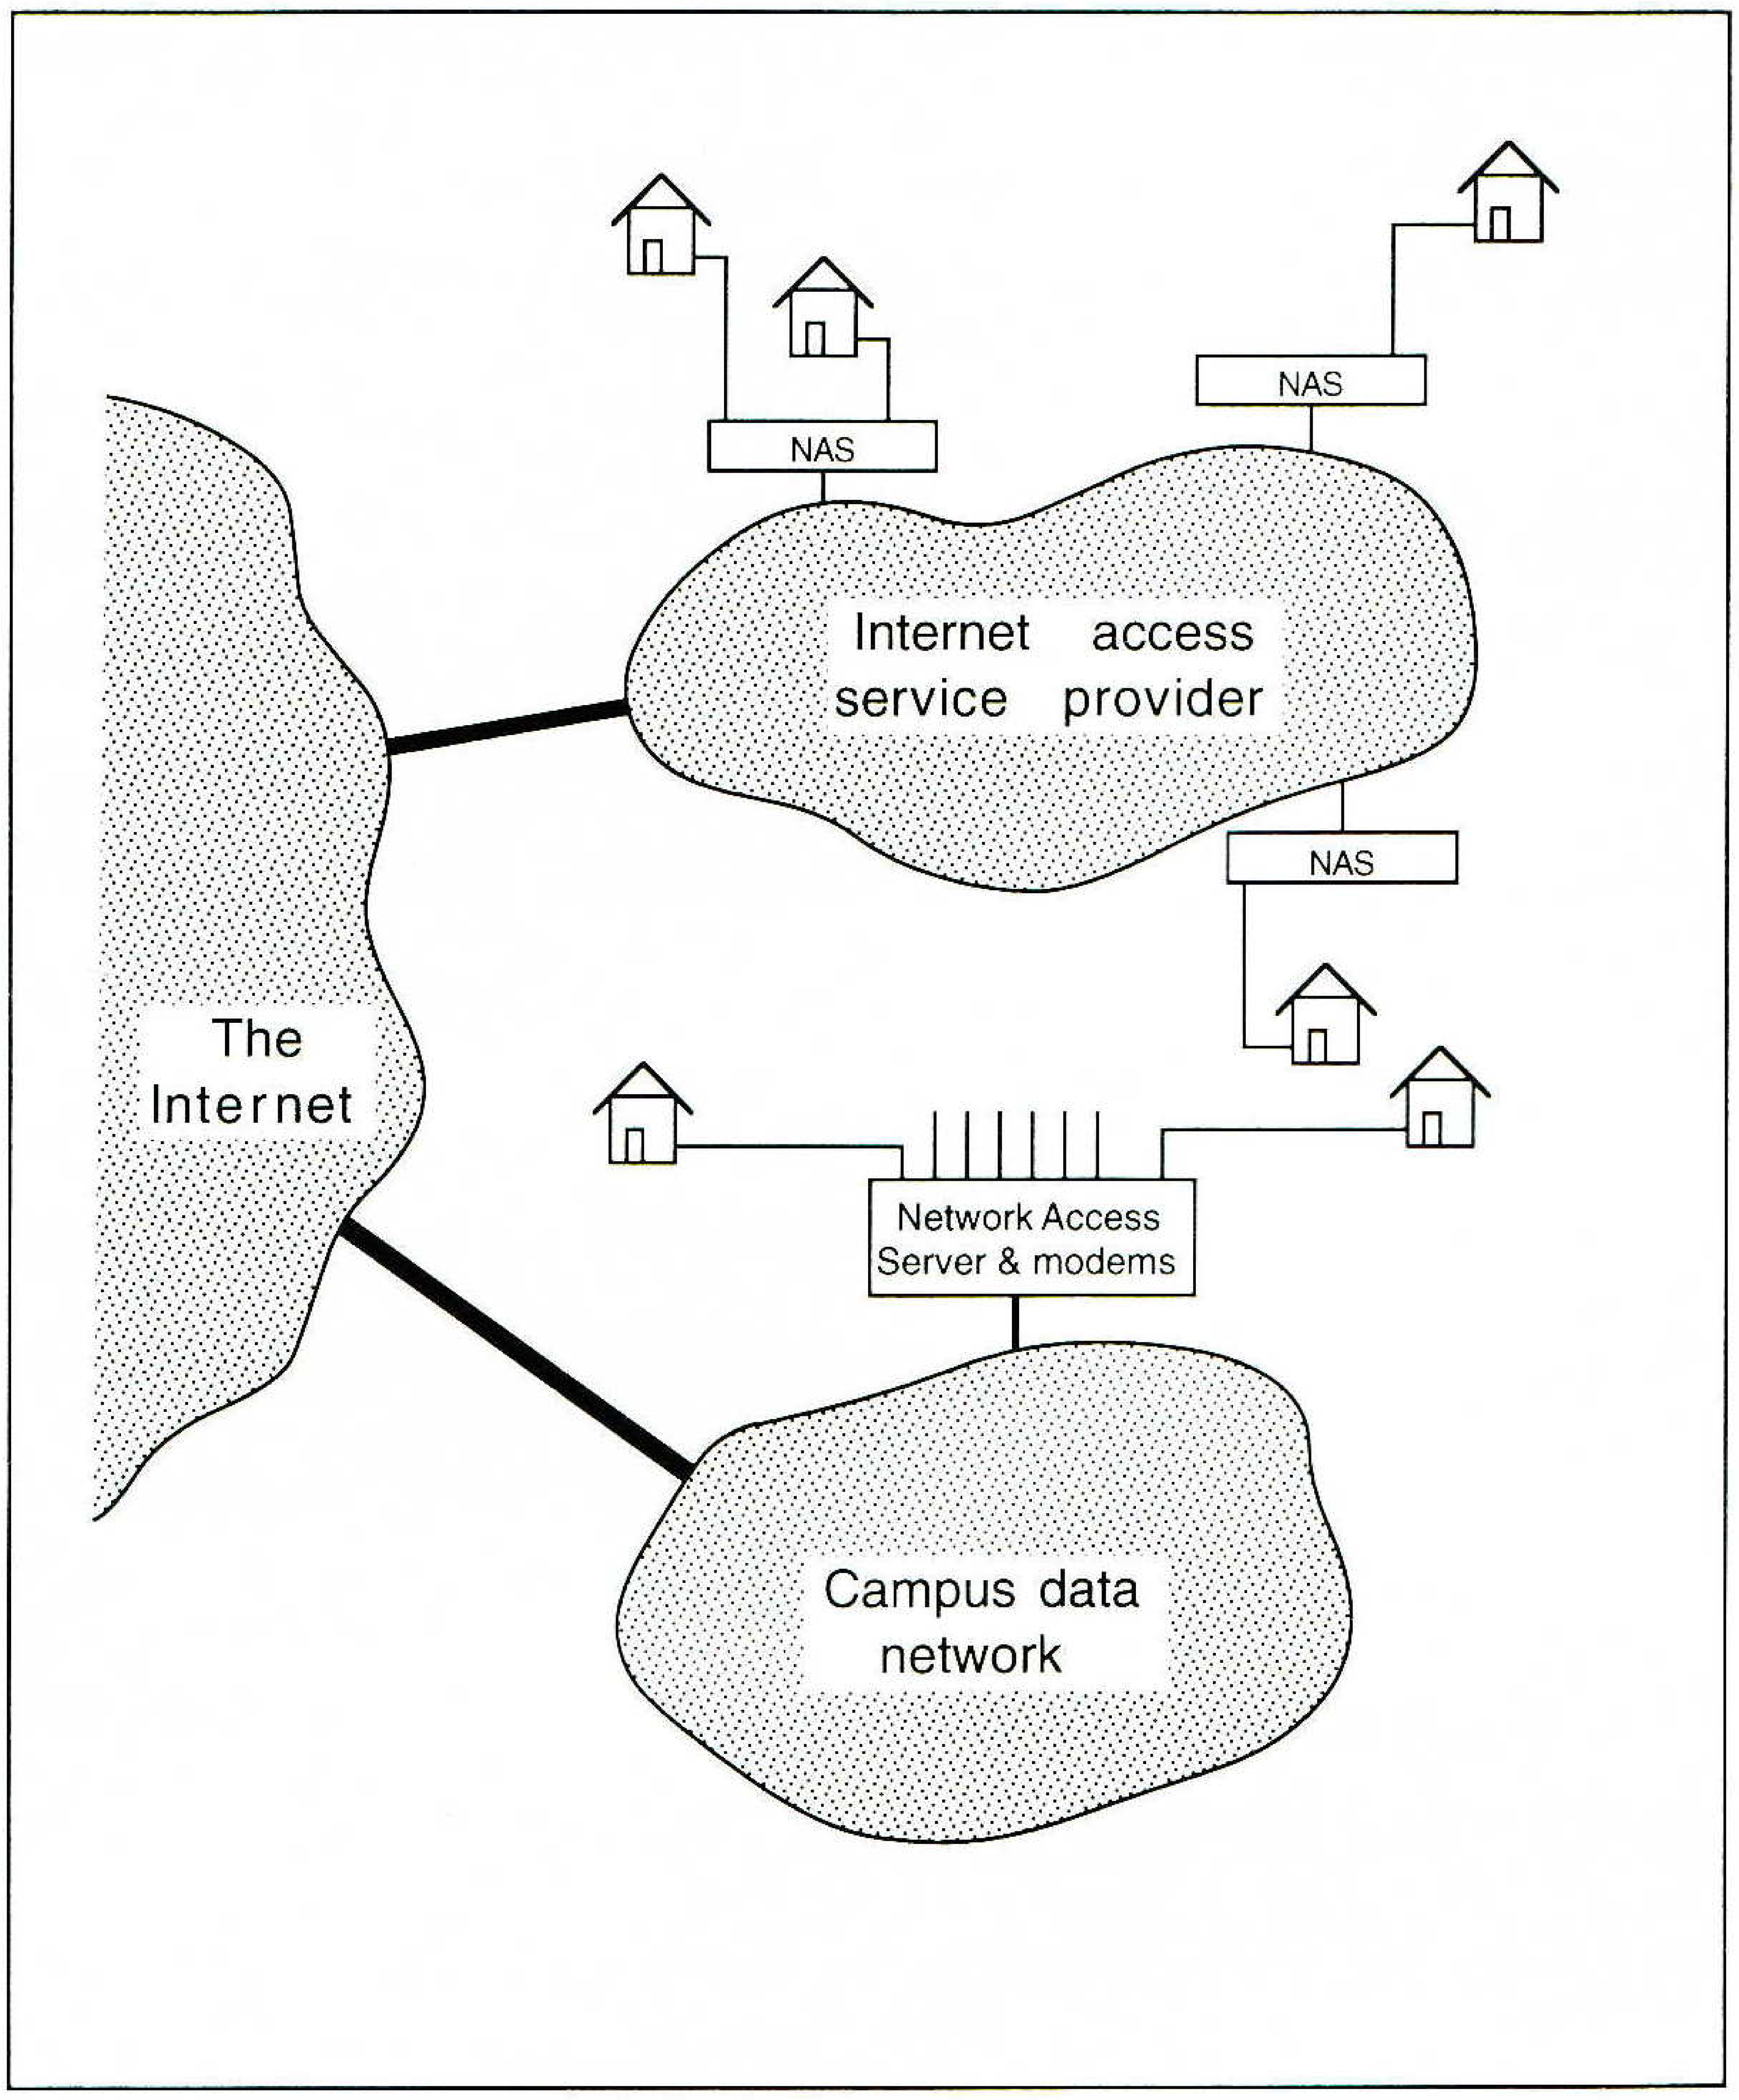

Figure 2: Commercial Alternative for Accessing the Campus Network

We see a full cost to the university of $60 to $80 per month to support a state-of-the-art modem in a modest sized service pool. Clearly this is too high to pass on to the average campus user directly, especially students. We certainly could adopt a cost model based on statistical "hold times" and prorate a given modem over a number of users. However, as users require more and more time "on line" and the peak use periods overlap, the reality of spreading the cost of a modem across a number of our high demand users is questionable. For the time being however, given a contention ratio of 15:1, a lower bound on monthly costs would be in the $5–10 per month range. Is this significantly less than commercial ISP services?

Commercial ISP accounts are offered in the $10–20 per month range today with a charging model that uses a combination of flat and measured service. A typical example is "$15 per month plus $2 per hour after the first 40 hours." While not nearly as inexpensive as the campus provided service on a "full time month basis," it is quite affordable in absolute dollars to the modest end user. Furthermore, many ISPs allow "free" access during off-peak times giving the determined student much more than the basic service time. It is reasonable to expect that the cost of such services will continue to come down as competition increases. It might be possible to negotiate a "bulk service agreement" with a commercial ISP to bring the cost down for university community members in return for campus guaranteed minimum number of accounts, rebilling of users, or user support functions.

Another concern for most universities is the level of services available from the ISPs. Most campus dial-in services support SLIP and PPP which makes the user's computer appear to be a node on the campus network. Currently many ISPs charge more for this type of service, in part because the support costs are higher and the market requiring this service is still smaller than that for simple "shell" access. University users require SLIP/PPP in order to make full use of information resources such as World-Wide Web and other client–server applications. An advantage of negotiating a bulk service contract with an ISP would be the ability to define the service to require SLIP/PPP support.

A related concern is how end users will gain access to campus restricted information. Many current systems or database servers that must restrict access to only "the campus community" use network addresses as the control mechanism. This is done because it is fairly simple to implement but it may not fully achieve either the campus's purpose or the end user's goals. If a user attaches directly to the campus network (i.e., with SLIP/PPP) they will have access to this information. If they attach to an ISP and then attempt to connect to the application, they will be seen as coming from a different network so may be denied access. Clearly the access control method is wrong, but there is no good alternative in wide spread use today. This is a problem that must be solved in any case because access control should be based on the user, not on where they are located on the network. A faculty member on sabbatical should be able to gain access to campus resources from a remote location; conversely a visiting scholar may not qualify to access campus licensed resources even though they happen to be "on the campus" temporarily.

Other concerns might include the adequacy of bandwidth between the ISP's modems and the campus network, the increased complexity of problem resolution, and a somewhat greater potential for eavesdropping on traffic. Privacy might also be an issue. As commerce increases on the Internet, ISPs might be tempted to augment revenues by selling information on subscribers preferences to interested marketing services. We must stay closely involved in public policy issues in addition to the purely technical service issues.

Reasons for using ISPs

There are many good reasons to use Internet Service Providers:

• *Toll free calling:* One great advantage that ISPs have that universities would find difficult to implement is "toll free" access for all users. (So-called 800 access is a very expensive option that we find quite unattractive.) In many areas, calls from beyond some distance (ca. 16 miles in the San Francisco Bay Area) are subject to measured service tolls. This creates a serious disadvantage for students, faculty and staff living at a distance from the campus. This disadvantage will be magnified as ISDN becomes deployed for Internet access.

The university could rent space in outlying areas and install phone lines, etc., but is this a reasonable activity for universities to engage in? ISPs typically will have or will create a service point wherever the density of subscribers warrants.

• *Access from anywhere:* A secondary advantage of using an ISP is to provide access while traveling. Many faculty and staff take portable computers to conferences and would like to retain access to campus resources. An ISP with nationwide coverage would allow this as easily as if the caller were local. Obviously the access control issues raised above may be a serious problem in this scenario.

continued on next page

Building the Urban Area Network *(continued)*

• *Competition and variety:* A commercial service that faces serious competition will be driven to provide innovative services and motivated to keep costs as low as possible. Newer communications technologies such as ISDN [5] and Frame Relay will be supported as soon as the market allows. Universities may be more limited in their ability to respond to new technologies because of their existing technology base and lack of an adequate funding stream to update it.

An obvious extension of the current dial-in connectivity model (connection of individual workstations) is connection of home networks. Building a LocalTalk or Ethernet network is fairly straight forward on a small scale. One computer could serve as a router to connect the entire network to the Internet with a single access line. This is a typical service model for an ISP. Are universities prepared to enter this service market?

• *The "apartment house" model:* We see the need to gain access to Internet services penetrating far beyond the university community. We have had inquiries from non-campus-owned living groups that would like to connect a building-wide network to the campus network. While many students live in the building, so do non-students. To connect the entire building seems beyond the scope of serving the university community. However, it fits perfectly the ISP model of connectivity. If we can solve the access control problem posed above, the ISP model could provide very good service to groups of students (and others) at relatively low cost through shared service. Apartment building managers would be motivated to add this service, (along with cable TV) in order to attract tenants.

• *Appropriate use of university resources:* Clearly an advantage of the ISP model is that most personal use would bypass the university's resource. Furthermore, when students graduate and no longer qualify as part of the "university community" they can retain their Internet identity and access more or less seamlessly. It could be retained for life.

The Urban Network Model

Where is all this leading?

One way to look at the problem is to ask "what would the ideal world look like in five years?" Just as voice telephone service has become ubiquitous, data network service will become essential to urban life. Learners of all ages will benefit greatly from access to the vast array of new information resources. Commerce will occur via this new pathway. Members of the "university community" will be distinguished not by who provides their "Internet dialtone service" but by their access to university research and learning resources via the network. However clearly we see this vision of the near future, it is just as clear that the university will not literally build the required urban network infrastructure.

As noted above Internet access has already become a commercial service. It is not a service the university should construct throughout the broader community any more than we construct voice dialtone services for the community. Today, the commercial network access services are embryonic and still rely on voice grade modems. We should be working as soon as possible with communications service providers to help define how such services will evolve by the end of this decade into a high capacity, affordable, ubiquitous "urban area network service."

Today the data communications service available to most homes is constrained by the physical characteristics of existing service drops and local loop cabling. We use increasingly complex modems over this media in order to try to get higher speed data transport between home and the network. ISDN is finally becoming available and will support at least 128 Kb/s to the home. However, for the types of applications we envision we will require data rates at the very least an order of magnitude higher. It seems unlikely that this capacity can be provided broadly and economically by current circuit-switched services, regardless of the media.

Existing high capacity services incur relatively high cost for the individual end user. T1, SMDS, or even Frame Relay have high installation and monthly costs and require very expensive termination equipment. Independent of cost, these point-to-point services would be hard to scale to the size required of an urban area. Consider existing voice service equipment sites as "hubs" trying to serve at least 50,000 subscribers at T1 with Frame Relay! We wouldn't build our campus area data networks that way.

Wireless data networks, including packet data services and nationwide data delivery services, are becoming popular. We believe this technology fills a particular need, but it seems likely that data rates will remain under 1 Mb/s for wide area services. (Microcellular services have demonstrated data rates in excess of 1 Mb/s, but have transmission ranges of less than 20 meters.) Also, since broadcast systems are in effect "shared media," the actual bandwidth available to any particular user may be quite constrained.

A very interesting and relatively new development may offer an opportunity to engineer a cost effective, scalable, high capacity packet data transport service. Both CATV and voice telecommunications service providers are planning to install new cable plants in major urban areas over the next decade. These likely will consist of "fiber to the neighborhood and coax to the home" (the so-called *Hybrid Fiber Coax* (HFC) cable plants [1]) with active electronics in neighborhood "hubs." This will support not only voice dialtone and the proverbial "500 channels of mud wrestling," but potentially megabit per second bi-directional transmission speeds for data services. What service model(s) might come as a result of this new infrastructure?

We believe that three key determinants of success will be: subscriber cost, scalability, and the ability to adapt to a broad potential market. One possible implementation model would be to view the new high speed infrastructure as a packet data transport cloud with information service providers accessible to anyone on the cloud. This "Packet Data Transport Service" (PDTS) could be quite similar to Frame Relay or SMDS. With the right engineering this PDTS could be easily scalable. Scalability and the fact that the cable plant cost would be recovered largely by the other services should keep the PDTS subscriber cost reasonable.

The PDTS service model would be very similar to that proposed for "video on demand" with information providers attaching to the new urban HFC infrastructure and selling services to any end user similarly connected. One generic type of data service provider would be the *Internet Access Provider* (IAP). IP packets would be relayed between the user and the IAP via the PDTS. Another major service provider easily could be the university: a provider of learning resources and research information. All that would be required would be a link between the campus network and the urban area network cloud. (See Figure 3)

continued on next page

Building the Urban Area Network *(continued)*

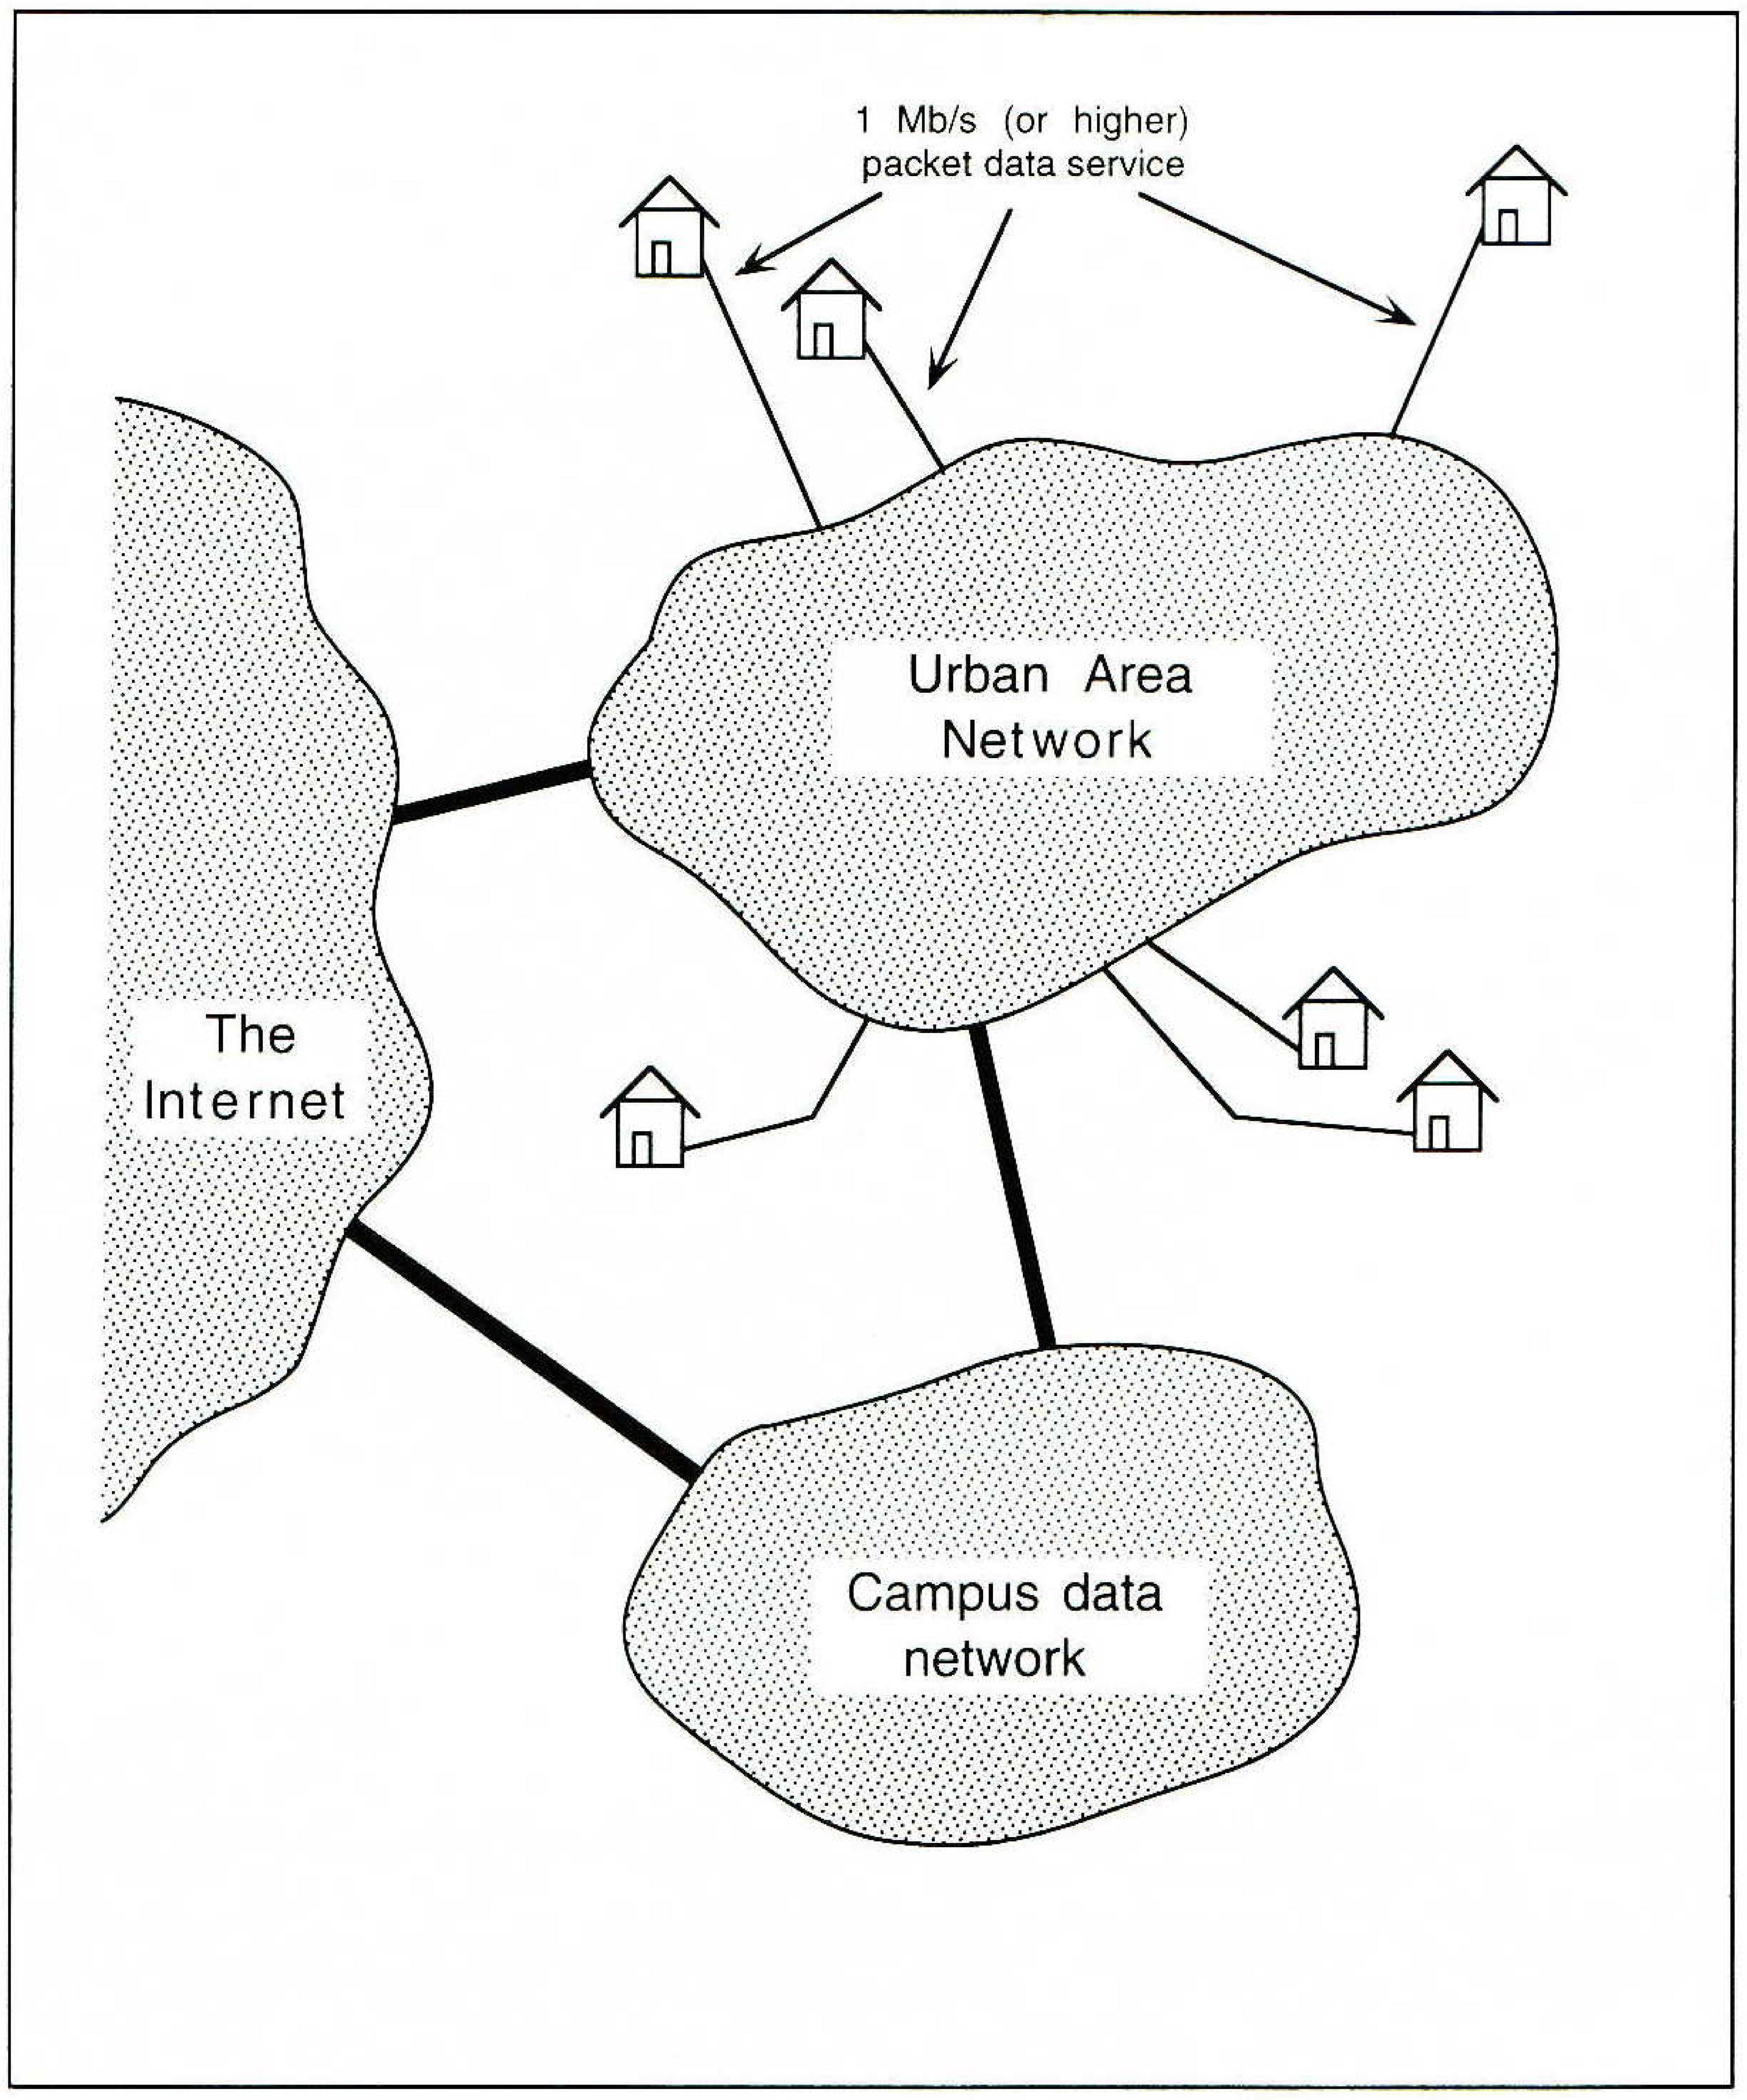

Fig 3: The University as a Service Provider on the Urban Area Network

The exact details of the technology that might provide high capacity "packet data dialtone" is a subject for another article. The points that become clearer by envisioning where we might want to be in five years are:

- The Urban Area Network (UAN) won't be built by the university,

and

- Use of the UAN will cost end users something.

If we accept this conclusion then adopting a strategy to outsource access to the campus network as soon as possible begins to get across the idea that this type of service in their homes will cost end users something—just like cable TV or voice dialtone. It also motivates us to solve the access control problems, and tells us we should be working closely in partnerships with communications technology developers to ensure that emerging technology will serve our needs.

How do we get there from here?

I suggest there are two main areas in which universities might have a significant and effective role. One is in developing robust access control standards and technology so that we can offer proprietary information resources within this new communications environment. The other is in the interface standards and the design and operating characteristics of urban area network technology.

The IETF and other standards bodies as well as commercial developers are working on standard security models for distributed authentication and secure data transport. These must be accompanied by robust authorization mechanisms as well as broadly available implementations for all popular platforms. These technologies will be critical for all network based applications.

The most critical area remains the UAN technology itself. Services being tried today that rely on traditional circuit switched access or shared access to low bandwidth data streams within standard CATV systems are unlikely to be able to carry the high volume of traffic we will require. There are engineering solutions that can work (for example the ATM over CATV work described in [1]) but we must take a lead role in convincing developers and service providers to try them.

By working closely with commercial communications service providers today to encourage wide spread deployment of urban network services, we will be fostering the very infrastructure that our campus communities will depend upon tomorrow.

References

[1] Laubach, M., "To Foster Residential Area Broadband Internet Technology—IP Datagrams Keep Going, and Going, and Going..," *ConneXions,* Volume 10, No. 2, February 1996.

[2] Lippis, N., & Herman, J., "Widening Your Internet Horizons: Wide-Area Options for Internets," *ConneXions,* Volume 5, No. 10, October 1991.

[3] Hobby, R., "The Point-to-Point Protocol (PPP). A new proposed standard Serial Line Protocol," *ConneXions,* Volume 4, No. 4, April 1991.

[4] Leifer, D., "ISDN: Why use it?," *ConneXions,* Volume 4, No. 10, October 1990.

[5] Leifer, D., "ISDN for Internet Access," *ConneXions,* Volume 10, No. 2, February 1996.

[6] Crowcroft, Jon and Handley, Mark "The World-Wide Web: How Servers Work," *ConneXions,* Volume 9, No. 2, February 1995.

[7] Berners-Lee, T., "A Summary of the WorldWideWeb System," *ConneXions,* Volume 6, No. 7, July 1992.

[8] Berners-Lee, T., R. Cailliau, A. Loutonen, H. F. Nielsen and A. Secret, "The World-Wide Web," *Communications of the ACM,* Volume 37, No. 8, August 1994.

[9] Crowcroft, Jon and Handley, Mark, "Problems with the World-Wide Web," *ConneXions,* Volume 9, No. 6, June 1995.

[10] Mark Handley and Jon Crowcroft, *The World-Wide Web: Beneath the Surf,* ISBN 1-85728-435-6, UCL Press, 1995.

[11] Lucien Rhodes, "The Race for More Bandwidth," (Interview with Milo Medin of @Home), *WIRED,* 4.01, January 1996, page 140.

DAVID WASLEY holds a Masters Degree from the University of California, Berkeley, and has been a member of the staff of the University for 27 years. For the last decade he has been responsible for the development of the UC Berkeley campus data network and associated services. He was active in the founding and development of the Bay Area Regional Research Network (BARRNet). He is co-author of RFC 1709 "K-12 Internetworking Guidelines." Recently he joined the UC Office of the President in order to focus on new issues and challenges in the area of Information Infrastructure Planning. E-mail: `David.Wasley@UCOP.EDU`

Opinion: A Statement Against the CDA

by John S. Quarterman,
Matrix Information and Directory Services

Introduction

For use in our lawsuit against the Communications Decency Act (CDA), my lawyer asked me for a personal and organizational bio, and "statements and concrete examples on how the CDA will affect your work, with a particular emphasis on how the CDA may also affect your print work." This is what I sent:

Statement

Many of my articles appear both on paper and online, often simultaneously. My books usually appear first on paper as books, but they usually draw at least partially from material I or others have previously distributed online, and the books themselves eventually appear in part or whole online. The subject matter treated ranges from extremely technical to sociological and personal. Must I, because of the CDA, now decide what to censor when I publish it online? Practically, this would mean the print versions would often also get censored in the same way, because my resources are not inexhaustible, and there are limits to the multiple variants I can produce.

MIDS is far from alone in publishing both online and in print. Many major print publishing houses now distribute at least samples of their work online. Many of them use the Internet to communicate with their authors in producing books. Even more of their authors use the Internet as a research tool. Publishers are thus exposed to the effects of the CDA even if they do not actually publish online at all.

I have extensive experience in collecting information in many countries throughout the world, and the nature of the local regime is always a consideration. Some things I am not willing to print for fear that the source of the information might suffer retribution. In the past such things have included the extent of networking in eastern Europe before the fall of the Soviet Union, the distribution of political newsgroups in certain Moslem countries, and access to certain decadent western publications (such as *The Economist*) in certain east Asian countries.

To this list must I now add any mention of online material produced or distributed in the United States and remotely related to sex, excretion, or abortion? What about killing children? I would most likely be referring to child processes in a highly technical discussion, but would some local D. A. know that? I've already seen Guardian Angels asking online about "kill files" so they could expose them. The only kill files I know about are text patterns indicating which USENET news articles to ignore, but will a judge or jury untrained in technical matters know that, or will they assume I'm keeping lists of snuff films? Has the U.S. joined the list of repressive governments that I have to watch out for? Rejoined them, actually. It was not the KGB that caused me not to write about USENET in eastern Europe: it was my *western* European informants' fear of the U.S. Commerce Department.

Much scientific research is published online these days, and much of that passes over the part of the Internet that is in the United States. Will researchers in France now have to watch what they publish about reproductive matters because any discussion of abortion may be prohibited by the CDA?

The CDA has already made the United States the laughingstock of Europe, where it is commonly viewed as the new Prohibition. Except that laughing may stop soon, since this U.S. precedent will encourage the European Union to pass the similar law it has been considering.

JOHN S. QUARTERMAN wrote the first book about the Internet and related networks, *The Matrix: Computer Networks and Conferencing Systems Worldwide,* Digital Press, 1990, and is currently working on its second edition. He is a co-author of five other books, *The Design and Implementation of the 4.3BSD UNIX Operating System,* 1989, *UNIX, POSIX, and Open Systems: The Open Standards Puzzle,* 1993, *Practical Internetworking with TCP/IP and UNIX,* 1993, *The Internet Connection: System Connectivity and Configuration,* 1994, and *The E-Mail Companion: Communicating Effectively via the Internet and Other Global Networks,* 1994, all from Addison-Wesley. He is Editor of the color *Matrix Maps Quarterly* and the monthly *Matrix News,* both about issues that cross network, geographic, and political boundaries, and both published by Matrix Information and Directory Services, Inc., (MIDS) of Austin, which also conducts demographic surveys and other research into the composition of the Internet and related networks and their users. He is a partner in Texas Internet Consulting (TIC), which consults in networks and open systems, with particular emphasis on TCP/IP networks, UNIX systems and standards. He is a partner in Zilker Internet Park, which provides Internet access from Austin, Texas. E-mail: `jsq@mids.org`

And every tinpot dictator in the world has now been dignified in restricting local press: after all, if the U.S., which preaches freedom of the press, sees no difficulty in stifling communications, why should they?

In my roles as consultant and Internet service provider (ISP), I gather experience to use as material for my writing. No real ISP can fail to include USENET news, that whipping boy of the pro-CDA forces. Most consulting clients want USENET news in-house, as well. Must I refuse to perform a major part of my work because Congress chooses to believe *Time Magazine* and the fraudulent research it published that claimed erroneously that most of USENET is pornography?

I have heard proponents of the CDA argue that this law merely brings the Internet to the level of other media. This is profoundly wrongheaded, because the Internet is not like other media. The Internet is not the telephone, it is not radio, it is not television; it is not a newspaper, magazine, or letter. The Internet includes services with features similar to all these traditional media, but it is not the same as any of them. It is much more basic. It is more like the printing press, or perhaps paper, or the electromagnetic spectrum. To confuse the Internet with the telephone (as Senator Exon and others have done repeatedly in the press and on the floor of the Senate) is like confusing a newspaper with the paper it is printed on, or confusing a pornographic movie with white light. They aren't even in the same category of things. You don't ban paper because somebody might scribble a dirty word on it; you don't ban light because somebody might use it to project a movie you don't like. Yet that is what the CDA is trying to do with the Internet.

The Internet is not even limited to communications; its primary historical function is actually resource sharing. I daily use computers in New York, Georgia, Texas, and sometimes California, the Netherlands, and other parts of the world as if I were simultaneously sitting in front of each of them, and I transfer data among them at will. I often write a program on a machine in one state, execute (excuse me: *run*) the program on a machine in another state, and display data from it on the screen of a machine in another state, then mail, FTP, or otherwise transfer some results to a colleague in another state. There is no essential difference to me between using a computer down the hall and one 2,000 miles away except the relatively minor considerations of connection speed and security. Except now some of that data might be actionable because of the CDA; I don't know. I don't even know what some of my source material says until I have it translated from whatever language I got it in, usually by sending it to somebody I know through the Internet. Who knows what some Italian or Thai or German might have written that may be perfectly acceptable where they are but is "indecent" in this great land where female nipples cannot be depicted in public but neo-Nazi demonstrations are legal? Must I examine every byte I transfer? According to whose local interpretation of the CDA?

Many companies use the Internet in a similar manner, tying together geographically separated offices to run their businesses. Large publishers, in particular, commonly do this. Must they now revert to sea mail because a couple of hundred U.S. politicians who have mostly never used the Internet and have no understanding of it have passed an absurd law?

Conclusion

The CDA directly threatens my livelihood; it threatens the business capabilities of my customers; and it threatens freedom of speech online and in print. It threatens these things both in this country and abroad. This law is wrong and should be overturned.

Announcement and Call for Papers

The Second *USENIX Workshop on Electronic Commerce* will be held November 18–20, 1996 at Claremont Hotel and Resort in Oakland, California. The workshop is sponsored by the USENIX Association and co-sponsored by Fisher Center for Information Technology Management, UC Berkeley, and the School of Information Management and Systems, UC Berkeley.

This workshop will provide a major opportunity for researchers, experimenters, and practitioners in this rapidly self-defining field to exchange ideas and present results of their work. It will set the technical agenda for work in the area of Electronic Commerce by examining urgent questions, discovering directions in which answers might be pursued, and revealing cross-connections that otherwise might go unnoticed.

Tutorials

The Workshop will begin with a day of tutorials. The tutorial program will offer a selection of tutorials from among several tracks on such topics as cryptography and security.

Workshop topics

Two days of technical sessions will follow the tutorials. Submissions are welcome for technical and position paper presentations, reports of work-in-progress, technology debates, and identification of new open problems. Birds-of-a-Feather sessions in the evenings and a keynote speaker will round out the program. We seek papers that will address a wide range of issues and ongoing developments, including, but not limited to:

- Advertising
- Auditability
- Copy protection
- Cryptographic security
- Digital money
- EDI
- Electronic wallets
- Hardware-enabled commerce
- Internet/WWW integration
- Legal and policy issues
- Negotiations
- Proposed systems
- Reliability
- Rights management
- Services vs. digital goods
- Smart-cards
- Anonymous transactions
- Business issues
- Credit/Debit/Cash models
- Customer service
- E-mail enabled business
- Electronic libraries
- Exception handling
- Identity verification
- Key management
- Micro-transactions
- Privacy
- Protocols
- Reports on existing systems
- Service guarantees
- Settlement

Questions regarding a topic's relevance to the workshop may be addressed to the program chair via e-mail to `tygar@cs.cmu.edu`. Proceedings of the workshop will be published by USENIX and will be provided free to technical session attendees; additional copies will be available for purchase from USENIX.

Submissions

Technical paper submissions and proposals for panels must be received by July 16, 1996. We welcome submissions of the following type:

- *Refereed Papers:* Full papers or extended abstracts should be 5 to 20 pages, not counting references and figures.

- *Panel proposals:* Proposals should be 3 to 7 pages, together with a list of names of potential panelists. If accepted, the proposer must secure the participation of panelists, and the proposer will be asked to prepare a 3 to 7 page summary of panel issues for inclusion in the Proceedings. This summary can include position statements by panel participants.

Please accompany each submission by a cover letter stating the paper title and authors along with the name of the person who will act as the contact to the program committee. Please include a surface mail address, daytime and evening phone number, and, if available, an e-mail address and fax number for the contact person. If all of the authors are students, please indicate that in the cover letter for award consideration. The program committee will offer awards of $500 for the best paper and the best student paper.

USENIX workshops, like most conferences and journals, require that papers not be submitted simultaneously to more than one conference or publication and that submitted papers not be previously or subsequently published elsewhere. Submissions accompanied by "non-disclosure agreement" forms are not acceptable and will be returned to the author(s) unread. All submissions are held in the highest confidentiality prior to publication in the Proceedings, both as a matter of policy and in accord with the U.S. Copyright Act of 1976.

Where to send papers

Please send submissions to the program committee via one of the following methods. All submissions will be acknowledged.

- *Preferred Method:* e-mail (*PostScript*): `ec96papers@usenix.org` Authors should ensure that their papers will print on a broad range of *PostScript* printers and submit in sufficient time to allow us to contact the author about alternative delivery mechanisms in the event of network failure.

- *Alternate Method:* 10 copies, via postal delivery to:
  Doug Tygar, Program Chair
  Computer Science Department
  Carnegie Mellon University
  5000 Forbes Avenue
  Pittsburgh, PA 15213–3891
  E-mail: `tygar@cs.cmu.edu` Fax: +1 412-268-5576

More information

If you have questions on the format of submissions or about the workshop, please call the USENIX Association office at +1 510 528-8649, or send e-mail to `ec96authors@usenix.org` or the program chair `tygar@cs.cmu.edu`. Materials containing all details of the technical and tutorial programs, registration fees and forms and hotel information will be available in September 1996. If you wish to receive the registration materials, please contact USENIX at:

USENIX Conference Office
22672 Lambert Street, Suite 613
Lake Forest, CA 92630
Phone: +1 714 588-8649 Fax: +1 714 588-9706
E-mail: `conference@usenix.org` URL: `http://www.usenix.org`

Or you can send e-mail to our mailserver at `info@usenix.org`. Your message should contain the line: send catalog. A catalog will be returned to you.

Important dates

Extended abstracts due:	July 16, 1996
Notification to authors:	August 5, 1996
Camera-ready final papers due:	October 7, 1996

Call for Papers

Multimedia Computing and Networking 1997 (MMCN97) will be held February 10–12, 1997 in San Jose, California. Advances in computer and networking technologies have fueled the rapid growth of research and development in multimedia computing and high-speed networking. As emerging multimedia technologies set higher performance levels at competitive costs, they are starting to enable and proliferate multimedia solutions in a spectrum of commercial and laboratory projects.

Topics

The objective of this conference is to bring together researchers, developers, and practitioners working in all facets of multimedia computing and networking. The conference will serve as a forum for the dissemination of state-of-the-art research, development, and implementations of multimedia systems, technologies, and applications. Presenters will be encouraged to make multimedia presentations and demonstrate their solutions.

Papers are solicited in all areas of multimedia, including, but not limited to:

- *Multimedia Computing:*
  - Hardware and software architectures
  - Multimedia operating system services
  - Real-time operating system services
  - Data streaming and delivery mechanisms
  - Media and user interaction
- *Multimedia Networking:*
  - Network and transport protocols
  - Quality-of-Service control and scheduling algorithms
  - Bandwidth management strategies
  - Synchronization mechanisms
  - Mobile network architecture
  - Community networking architecture
- *Multimedia and the Internet:*
  - Imaging on the Internet
  - Internet appliances
  - Intelligent network applications
  - Network programming languages
- *Multimedia Applications:*
  - Video-on-demand servers and services
  - Web servers and services
  - Search engines
  - Set-top technologies and operating systems
  - Multimedia conferencing and mail
  - Education and training
  - Digital libraries
  - Medical applications
  - Cyberspace communication, presentation, and interaction
  - Electronic communities
  - Entertainment and games
- *Multimedia User Interfaces and Authoring Systems:*
  - Video widgets
  - Synthetic animation
  - Intelligent information access
  - Interactive navigation schemes
  - Scripting languages
  - Authoring metaphors and editing techniques

Submissions

Please submit full papers for review. The submissions should not exceed 15 single-spaced pages including figures, tables, and references, using a typeface no smaller than 10 points. To expedite the reviewing process, please e-mail the paper in *PostScript* format to `mmcn@cs.utexas.edu`. Additionally, please send 1 hard copy to:

Professor Harrick M. Vin
Department of Computer Sciences
Taylor Hall 2.124
The University of Texas at Austin
Austin, TX 78712–1188

In addition, please submit electronically (in plain text format) a cover page to `mmcn@cs.utexas.edu`. Each cover page should contain: Title of paper, author names and affiliations, name and address (both postal and electronic) of contact author, abstract (500 words), keywords, and submission area (from the list of areas in the Call for Papers).

Publication

Each paper will be reviewed by the members of the program committee. Authors of accepted papers will be asked to submit a camera-ready manuscript that will appear in the conference proceedings. The Conference Chairs and Program Committee will also ask authors of the best papers to significantly enhance their papers and make journal form submissions to ACM/Springer Verlag *Multimedia Systems Journal*. A special issue of this journal will be devoted to the theme of the conference. Similarly, authors of selected papers will be asked to submit a tutorial style paper for *IEEE Multimedia Magazine*.

Important dates

Electronic submission deadline:	July 16, 1996 (Hard Deadline)
Deadline for receiving a hardcopy:	July 19, 1996
Notification of acceptance:	September 30, 1996
Camera-ready papers due:	December 16, 1996 (Hard Dln.)

More information

The call for papers as well as the deadline information can also be obtained from `http://www.cs.utexas.edu/users/mmcn`

Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use it for letters to the Editor, requests for the index of back issues, questions about particular articles etc.:

ConneXions—The Interoperability Report
303 Vintage Park Drive
Foster City, California 94404–1138, USA
Phone: +1 415-578-6900
Internet: `connexions@interop.com` `http://www.interop.com`

Subscription information

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 610-892-1959 outside the USA. This is the number for our subscription agency, Seybold Publications. Their fax number is +1 610-565-1858. The mailing address for subscription payments is: P.O. Box 976, Media, PA 19063–0976.

CONNEXIONS
303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED